The Latest News in DB2 Application Development

Bill Bireley

IBM Corporation

bireley@us.ibm.com

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

- Application Development Overview, Market Forces
  - Mobile and Cloud
    - Worklight
    - z/OS Connect
  - Big Data and Analytics
    - IDAA Update
    - DB2 BigInsights Connector
  - Agile Development and DevOps
    - DB2 JSON API
    - RDz
- DB2 11 SQL and Application Update
- Language Interface and DB2 Connect Update

Enterprise innovation will be realized through the integration of new era technologies with core systems

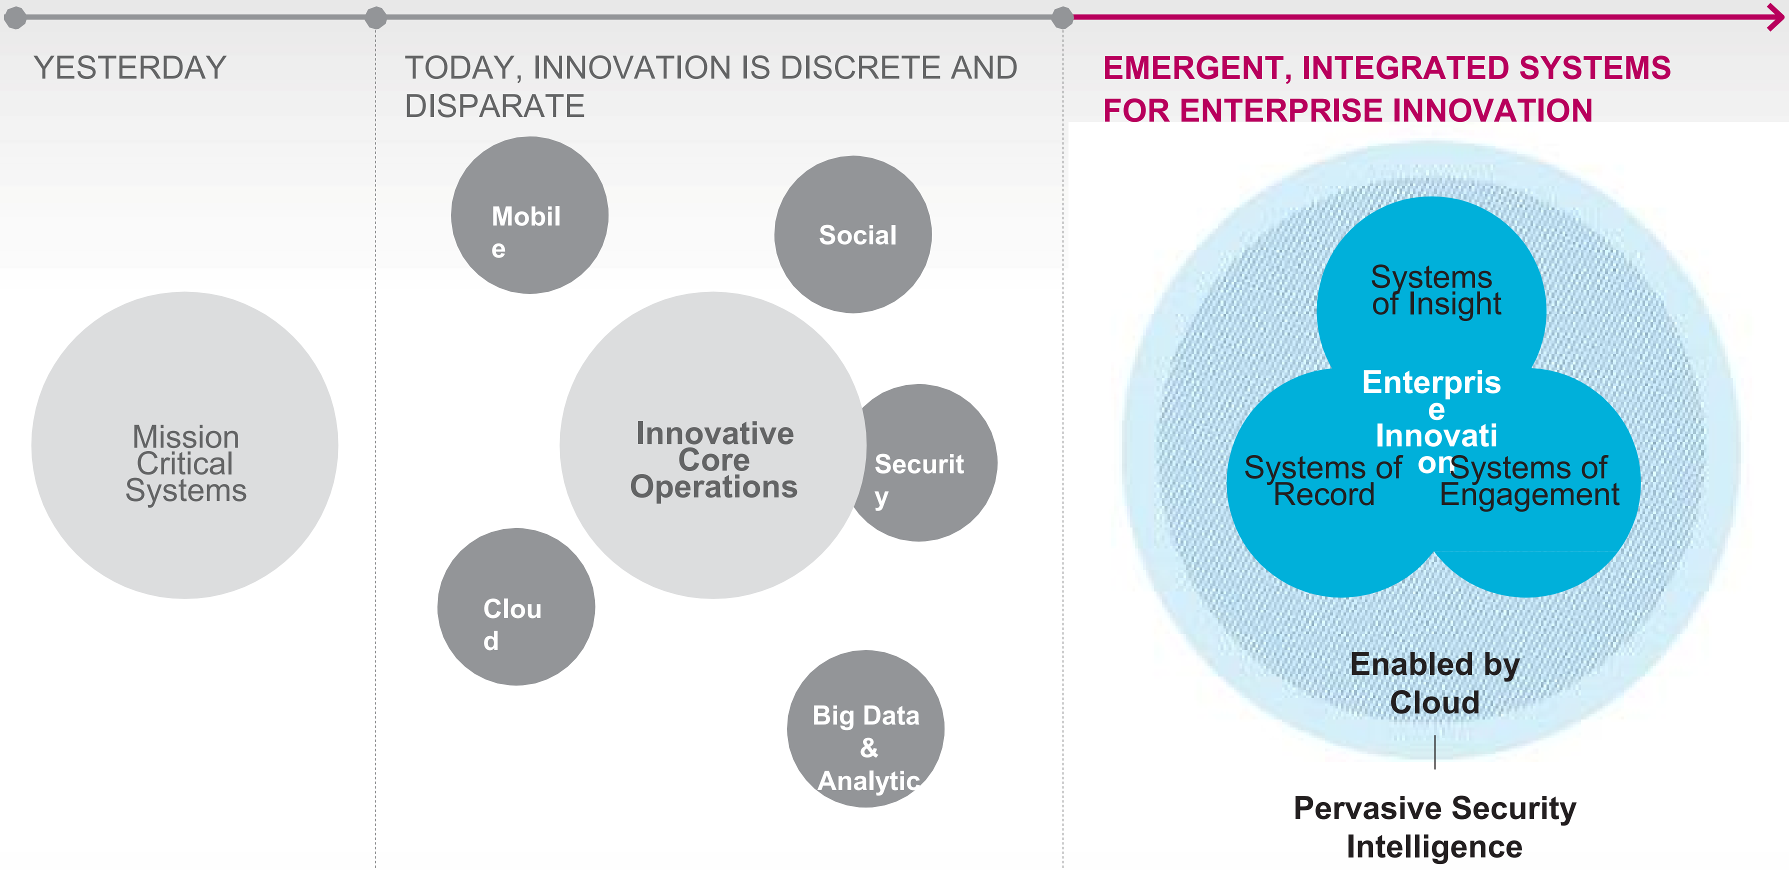

Leaders will take a systematic and integrated approach to drive enterprise innovation

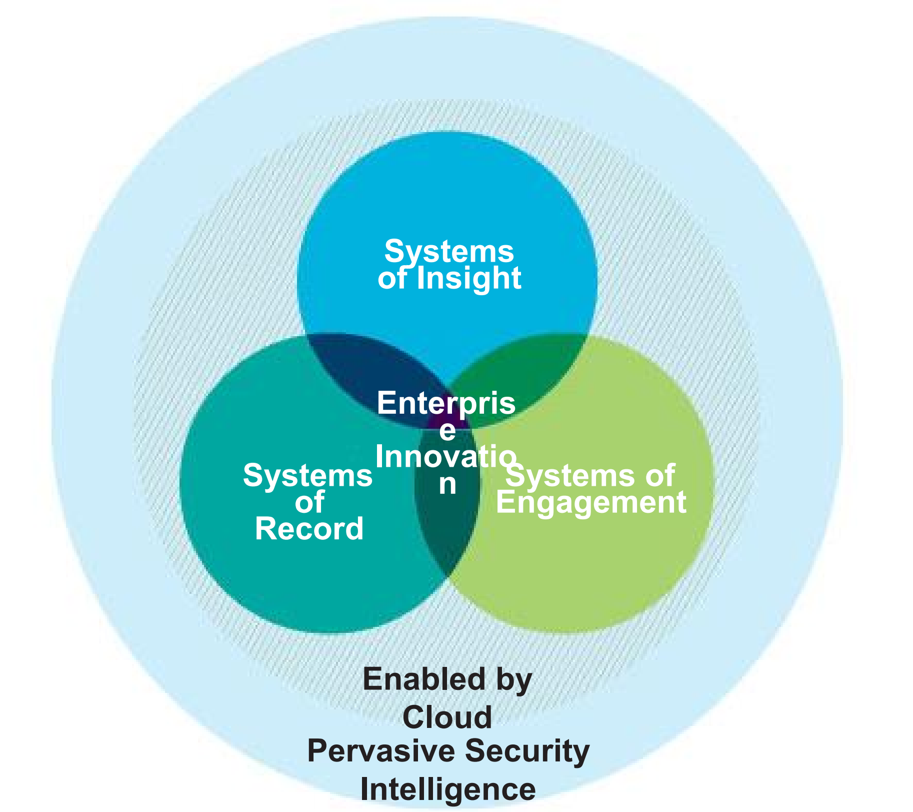

IBM works with you to build and integrate:

Systems of Engagement

How are you harnessing mobile and social to transform relationships with customers, employees and citizens?

Systems of Insight

How are you harnessing Big Data for competitive advantage?

Systems of Record

How are you transforming your existing IT systems for the requirements of the new era?

System z in a Mobile World

The role of System z in your mobile strategy

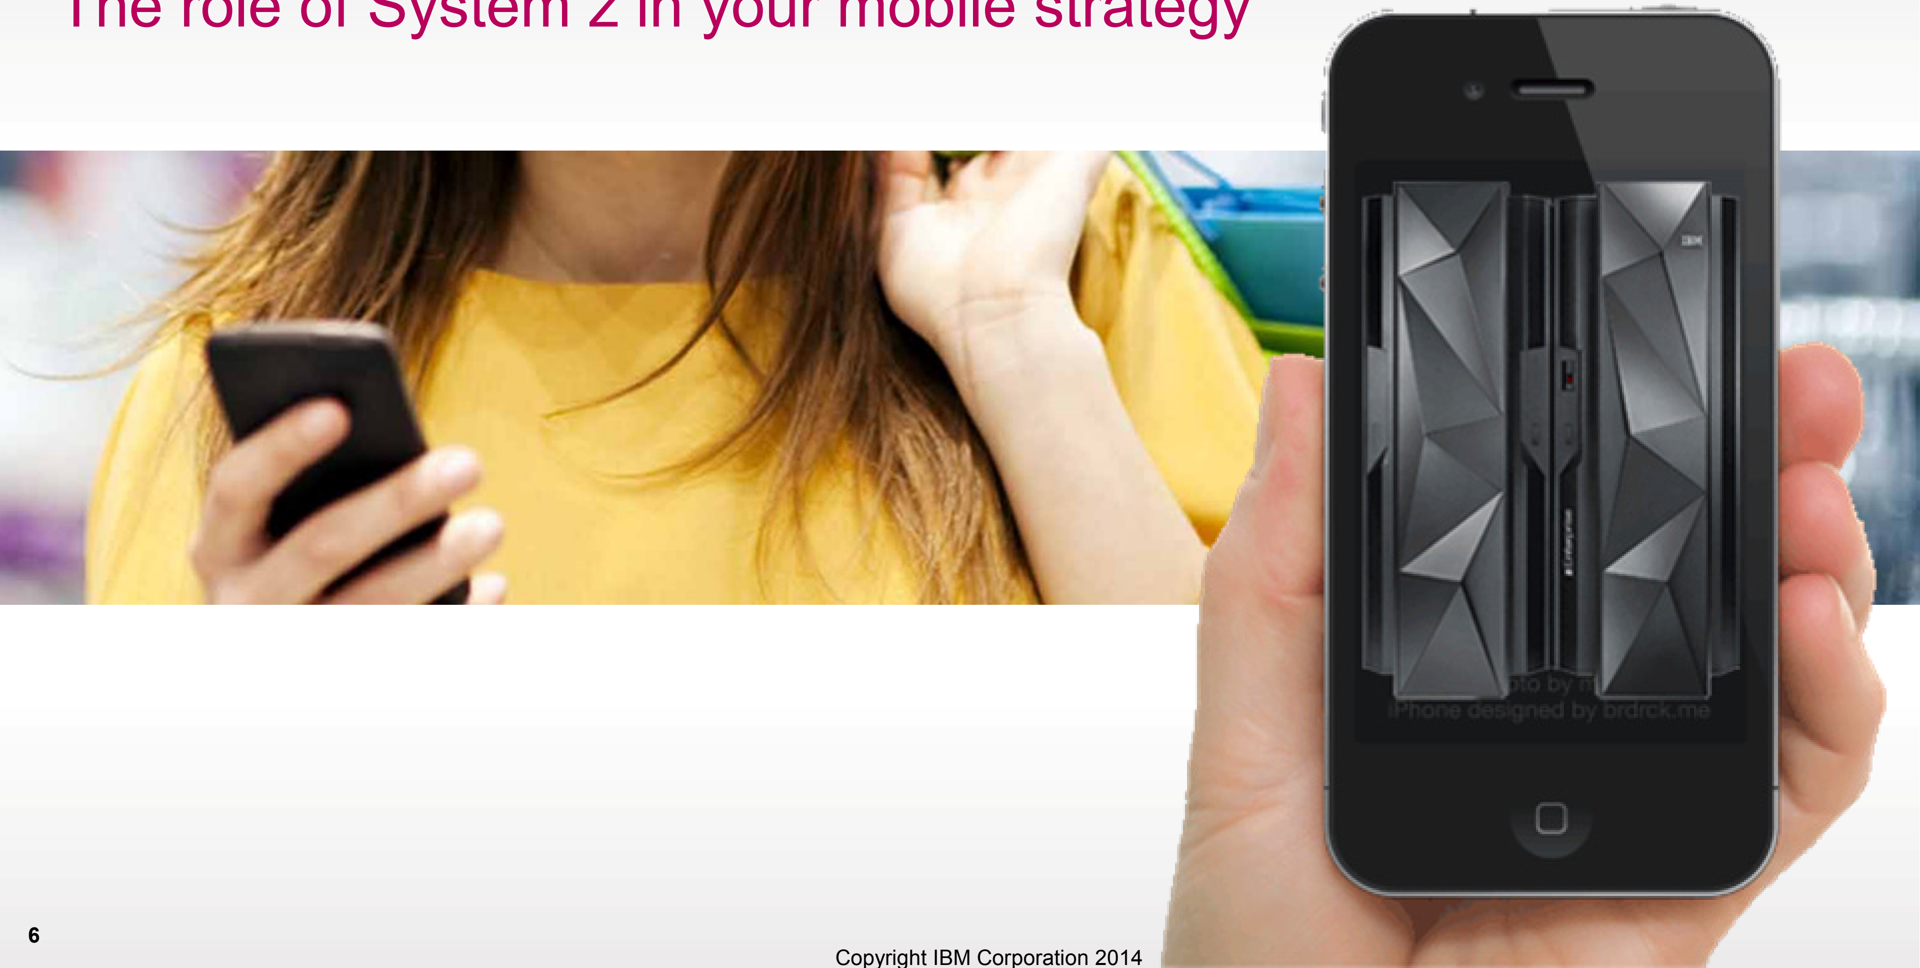

System z Unique Characteristics to support Mobile Applications

- **Easy-to-consume APIs** from CICS, DB2, IMS allow you to leverage your investment in z/OS transactions to quickly add a mobile channel.
- z/OS enables **massive and simple scalability** in a single footprint, to handle the workload of millions of devices and sensors
- Worklight **security** integrates with z/OS security providing end-to-end security and data privacy for mobile apps.
- z/OS **Workload Management** ensures your crucial applications remain responsive during sharp spikes in demand.
- **Low-latency** access to critical enterprise data. Mobile usage patterns favor short, read-only data requests (Users check account balances) So fast access to operational data, with low latency, is key.
- Business Resiliency for critical mobile apps

"70% of mobile applications touch a mainframe system"
Infrastructure matters for mobile applications. The System z platform's scalability, security, and resilience can enhance critical mobile applications.

Copyright IBM Corporation 2014

Enabling the Composable Business for access to z/OS assets

IBM WebSphere Liberty z/OS Connect

- **Simplify mobile and cloud connectivity to z/OS applications and data**
- **Capture new opportunities from existing enterprise investments**
- **Provide auditing and chargeback for cloud, mobile and web based requests**
- **Reduce development time and cost via industry standard interface to enterprise assets**

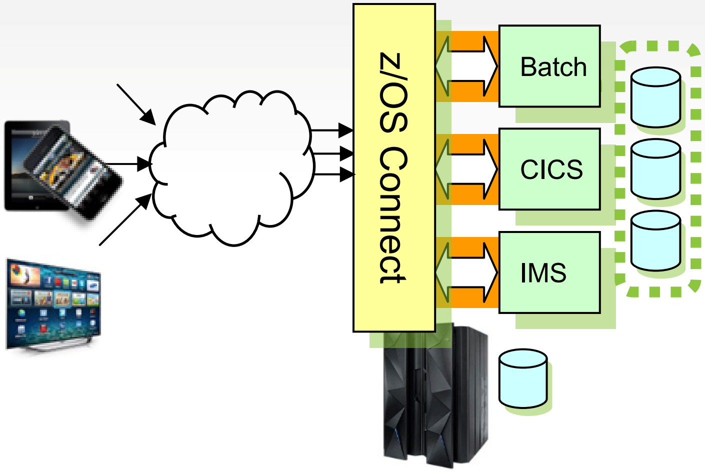

What's new in IBM WebSphere Liberty z/OS Connect

✓ ***Secure REST based connectivity to enterprise applications and data***

✓ ***Tight integration with System z and z/OS qualities of service***

✓ ***Delivered as a feature of the Liberty Profile in WebSphere, CICS and IMS***

System z addresses Enterprise mobile development and delivery challenges

Fragmentation and developing for multiple mobile platforms

- Highly fragmented set of devices, platforms, languages, and tools complicates development, test, and operations

IBM Worklight Studio and RDz

- Seamless integration with z data and transactions.
- Device runtime provides mobile device independence.

Accelerated time to market requirements

- Accelerated development demands instant provisioning of development servers.
- Spikey mobile traffic demands highly scalable cloud-based infrastructures, for both SoE and SoR.

System z Scalability

- System z Linux cloud enables rapid provisioning of Worklight servers.
- z/OS is the leader in transaction processing and easily scalable to handle workload increases.

Connecting apps with existing enterprise systems

- Apps typically need to leverage existing enterprise services, which must be made mobile-consumable, and remain secure.
- Enterprise systems must be able to instantly provision new services and environments.

z/OS is mobile enabled

- z/OS subsystems are mobile-ready, with consumability enhancements planned. Eg: MongoDB API for DB2
- End to end mobile security.
- High-performance access from z/Linux

IBM Worklight overview

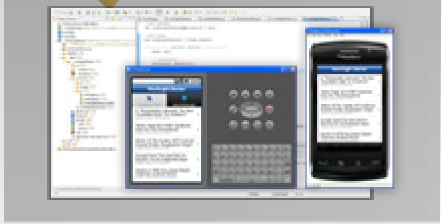

Worklight Studio

The most complete, extensible environment with maximum code reuse and per-device optimization

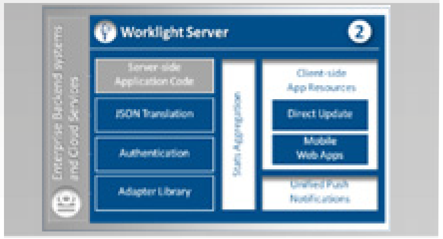
Worklight Server
2
Enterprise Backend systems and Cloud Services
Server-side Application Code
JSON Translation
Authentication
Adapter Library
Stats Aggregation
Client-side App Resources
Direct Update
Mobile Web Apps
Unified Push Notifications

Worklight Server

Unified notifications, runtime skins, version management, security, integration and delivery

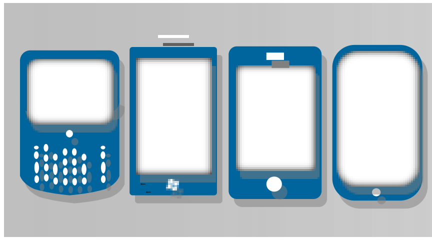

Worklight Device Runtime Components

Extensive libraries and client APIs that expose and interface with native device functionality

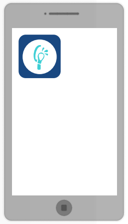

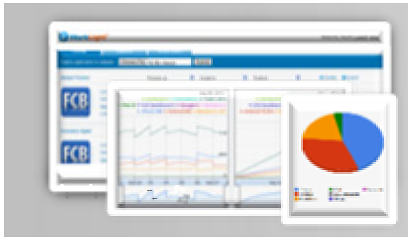

Worklight Console

A web-based console for real-time analytics and control of your mobile apps and infrastructure

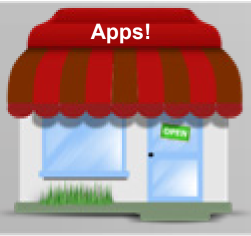

Worklight Application Center

Set up an enterprise app store that manages the distribution of production-ready mobile apps

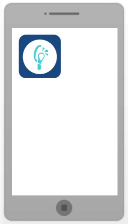

Benefits of developing and deploying with Worklight

- Worklight Studio:
  - Supports all common mobile device types
    - Can customize for each type using “skins” via the server
  - Very flexible choice of design patterns:
    - Pure HTML5
    - Use compatible Javascript framework on device – (jQuery, Dojo…)
    - Use Apache Cordova to access device functionality (GPS, compass, camera…)

- Worklight Server
  - Support hundreds of thousands of users
  - Directly update device applications
  - Hierarchical to JSON data conversion
  - Push notification to users
  - Security Framework – control access to protected resources
    - Robust mechanisms for client to server authentication, integration with back-end service security

IBM z/OS Connect

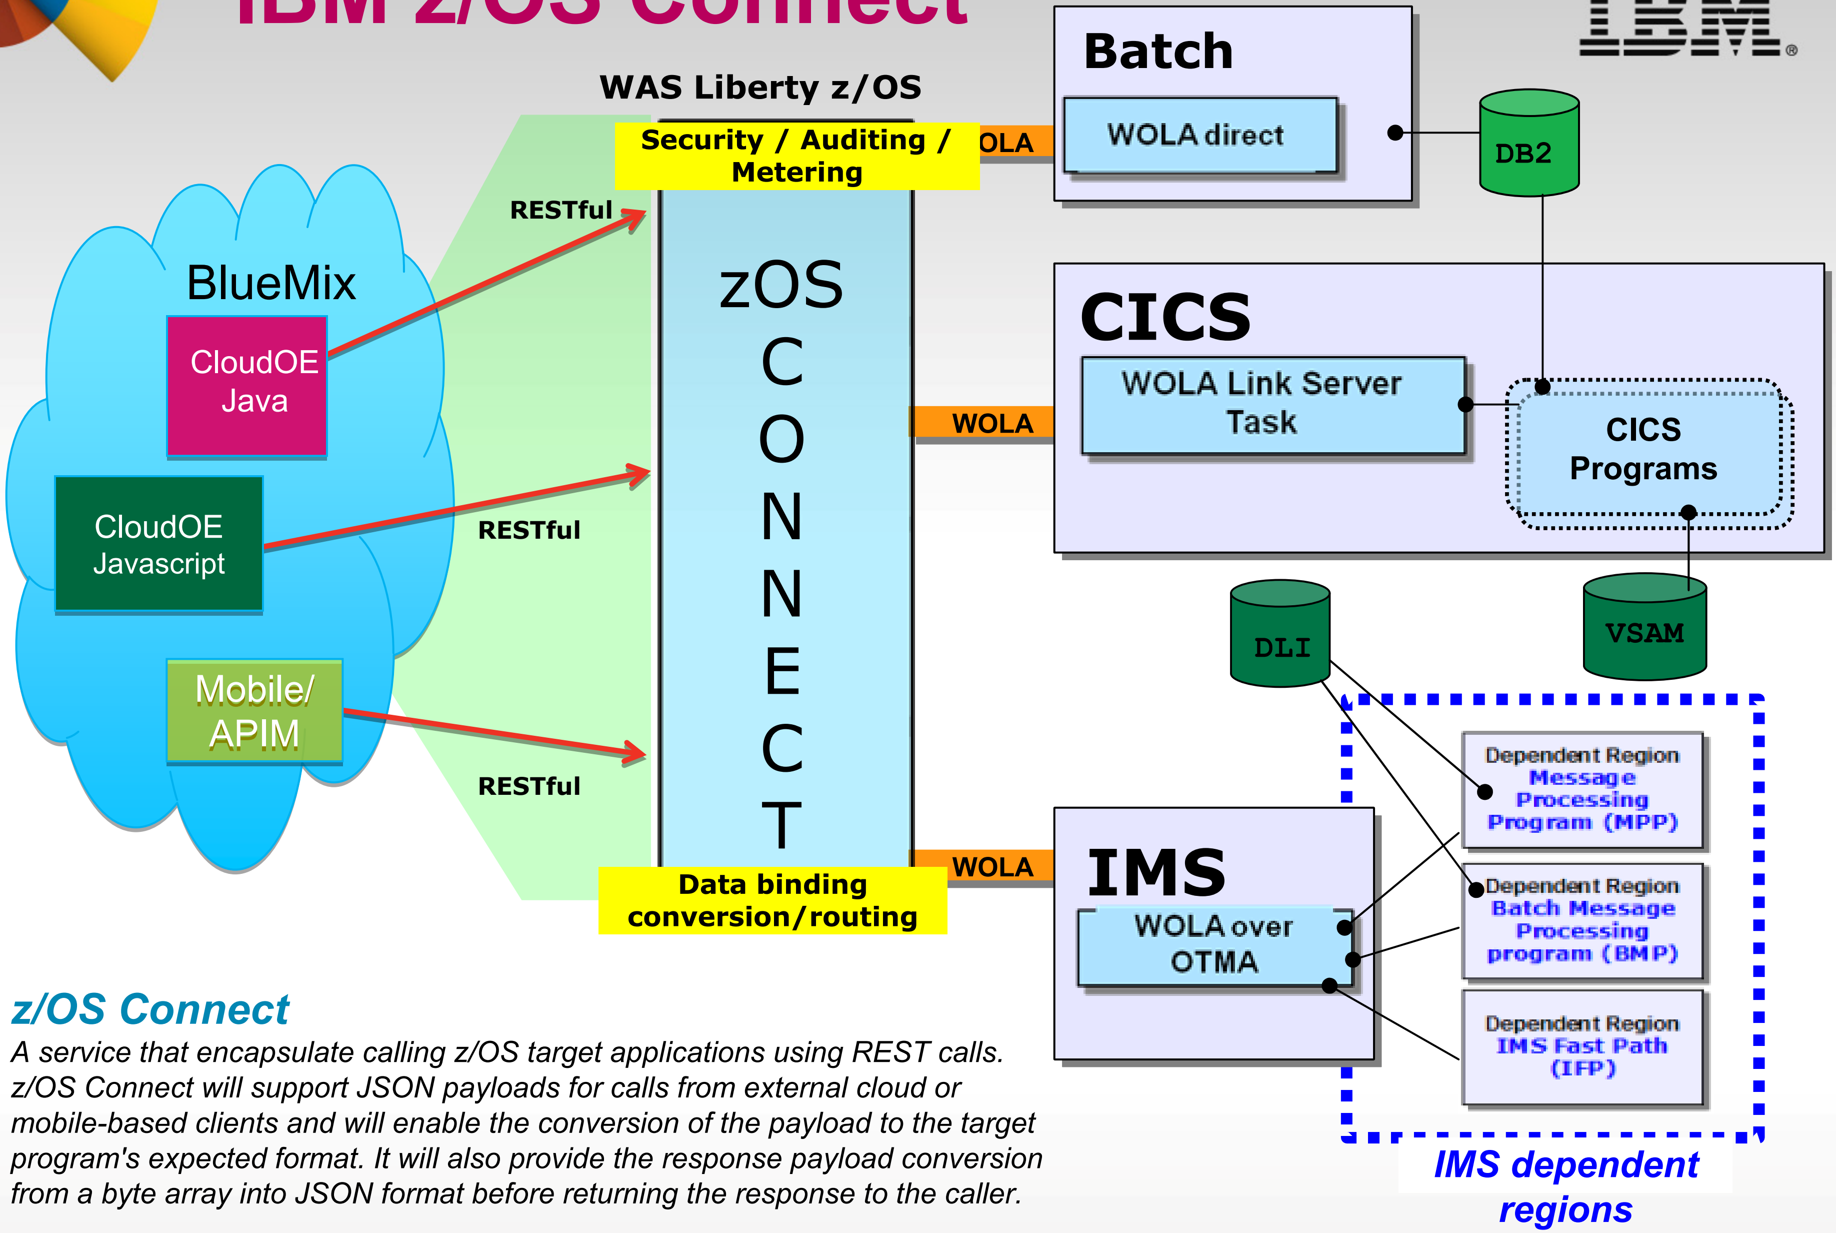

z/OS Connect

A service that encapsulate calling z/OS target applications using REST calls. z/OS Connect will support JSON payloads for calls from external cloud or mobile-based clients and will enable the conversion of the payload to the target program's expected format. It will also provide the response payload conversion from a byte array into JSON format before returning the response to the caller.

Big Data and Analytics

What is Big Data? A few definitions

- “Big data is the term for a collection of data sets **so large and complex that it becomes difficult to process using on-hand database management tools** or traditional data processing applications.” **- Wikipedia**

- "Big data is high volume, high velocity, and/or high variety information assets that **require new forms of processing** to enable enhanced decision making, insight discovery and process optimization." **- Gartner**

- “Big data is a collection of data from traditional and **digital sources inside and outside your company that represents a source for ongoing discovery and analysis**.” **- Lisa Arthur, Forbes Article**

Big Data and Analytics Options

Integrate DB2 OLTP data and workloads

- Some Big Data Options when Transaction Data is on DB2 for z/OS:
  - – In-database ETL (“light” analytics)
  - – ETL offload to Cognos, BLU, Infosphere Warehouse, etc.
  - – Data Analytics Accelerator (IDAA) – Integrated HTAP solution
  - – DB2 BigInsights (Hadoop) Connector
  - – 3rd party vendor integration – e.g. Veristorm and others
    - Run hadoop on System z (or other Linux systems)

Recent IDAA Programming Enhancements

- Version 3:
  - Incremental update – faster, less resource intensive refreshes
  - More query routing control
  - Support for additional types of queries that can be accelerated
  - Support for DB2 OLAP functions
  - Support for multiple-byte EBCDIC encoding
  - Support for additional DB2 scalar functions
  - Support for DECFLOAT for implicit casting
- Version 4:
  - Static SQL (the most requested feature of the Accelerator)
  - Multiple-row fetch for retrieving large answer sets from the Accelerator
  - Implicit casting when comparing VARCHAR and numeric data
  - Additional scalar function support (bringing the total number of scalar functions supported to over 80)
  - Support for EBCIDC and Unicode encoding schemes in the same DB2 subsystem and accelerator

BigInsights “Connector” for DB2 for z/OS

What is hadoop ?

- An open source software framework that supports data-intensive distributed applications
  - High throughput, batch processing
  - runs on large clusters of commodity hardware
    - Yahoo runs a 4000 nodes Hadoop cluster in 2008
- Two main components
  - Hadoop distributed file system
    - self-healing, high-bandwidth clustered storage
  - MapReduce engine

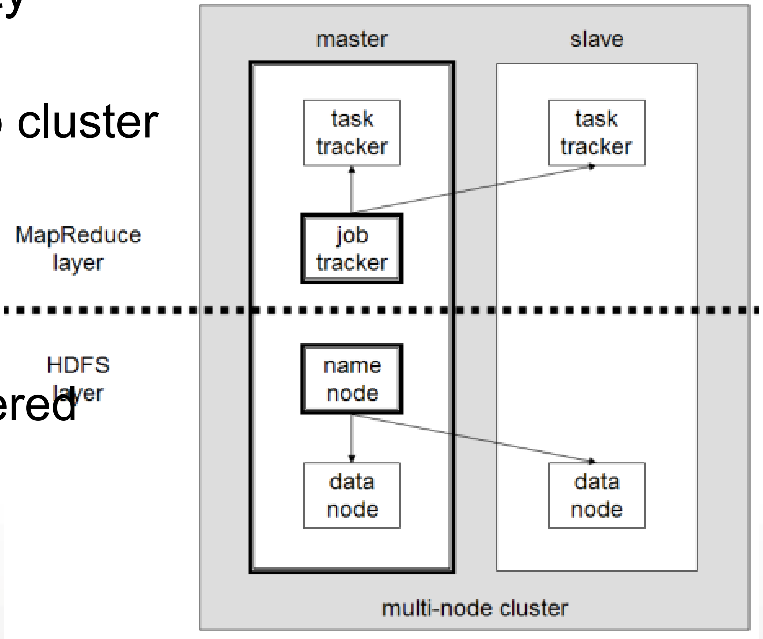

Hadoop: The underlying principle

- Lots of redundant disks – really inexpensive disks
- Lots of cores – inexpensive cores working all the time
- Disks crash – that's ok – just replace them
- Processors fail – that's ok – just replace them
- Network errors happen – that's ok - just retry
- Disks, processors, are all networked - the software manages data duplication, failover, restoration

Enhancing DB2 Analytics on “z” with Big Data

- DB2 is providing the connectors and the DB capability to allow DB2 applications to easily and efficiently access data in Hadoop

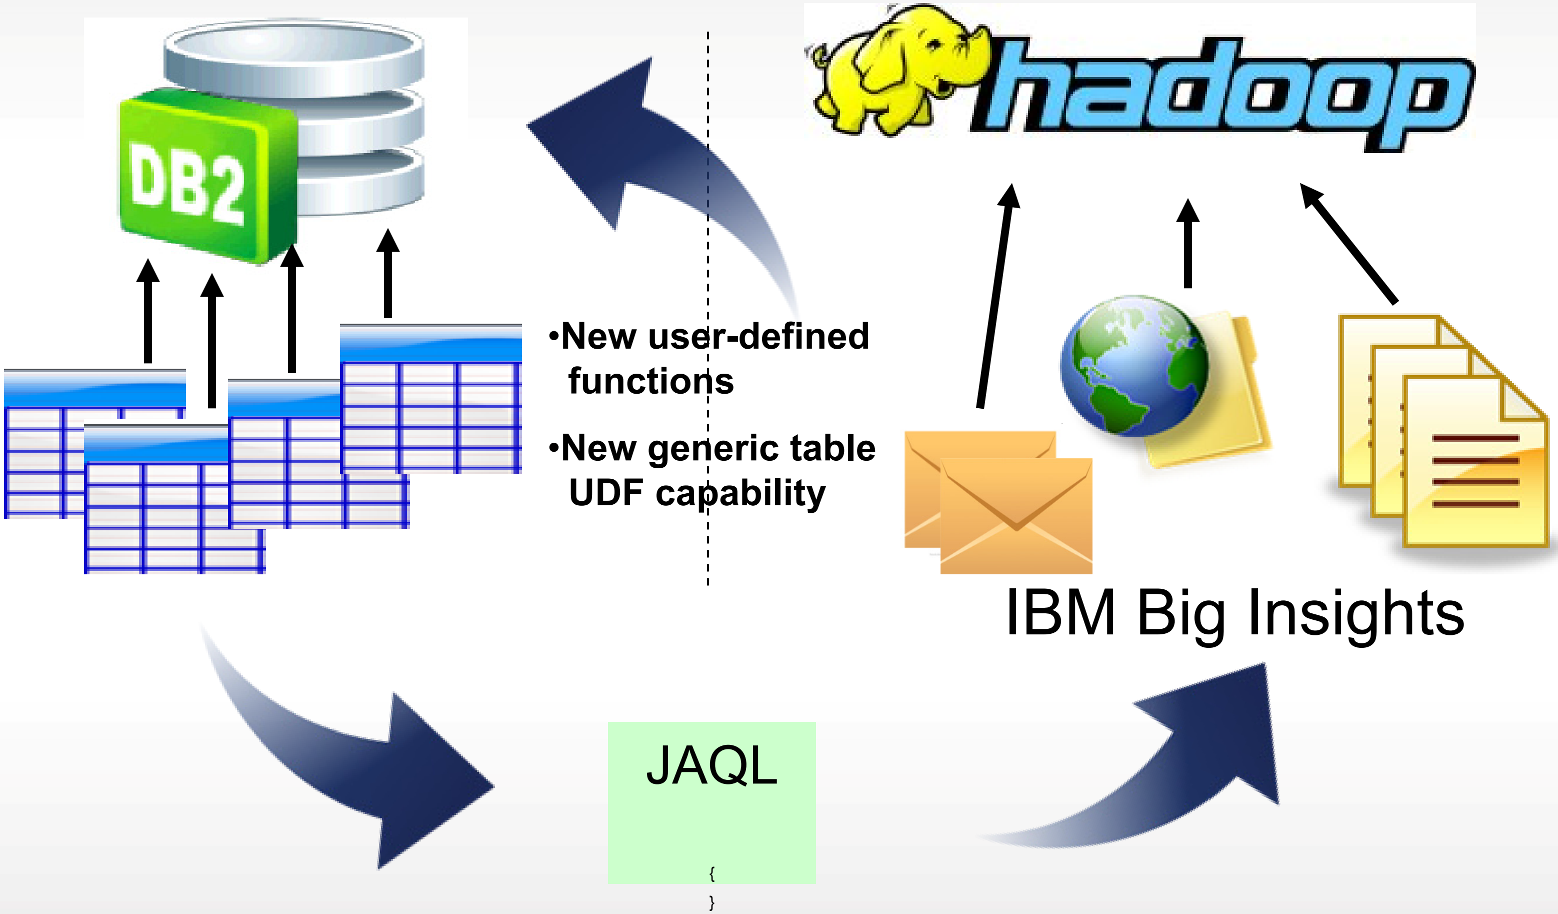

- New user-defined functions
- New generic table UDF capability

IBM Big Insights

JAQL

{

}

BigInsights “Connector” for DB2 for z/OS

- Two DB2 “sample” functions:
  - JAQL_SUBMIT – Submit a JAQL script for execution on BigInsights from DB2
  - HDFS_READ – Read HDFS files into DB2 as a table for use in SQL

- Notes:
  - Functions are developed by DB2 for z/OS
    - Shipped with DB2 11 in *prefix*.SDSNLOD2
    - Functions are *not* installed by default
  - Functions and examples are documented by BigInsights
    - http://www.ibm.com/support/docview.wss?uid=swg27040438

JAQL_SUBMIT Example

JAQL_SUBMIT

Submit a JAQL script for execution on BigInsights from DB2

```
SET RESULTFILE =
 JAQL_SUBMIT
   ('syslog = lines("hdfs:///idz1470/syslog3sec.txt");

     [read(syslog)->count()]->
      filter(strPos($,"$HASP373")>=0)->
      write(del(location="hdfs:///idz1470/iod00s/lab3e1.csv"))


'http://bootcamp55.democentral.ibm.com:14000/webhdfs/v1/idz1470/
    iod00s/lab3e1.csv?op=OPEN',


    'http://bootcamp55.democentral.ibm.com:8080',

    ''
   );
```

JAQL script containing the analysis

Intended HDFS file to hold the result

URL of the BigInsights cluster

options

JAQL – JSON Query Language

- Java MapReduce is the "***assembly language"*** of Hadoop.
- JAQL is a high-level query language included in BigInsights with three objectives:
  - Semi-structured analytics: analyze and manipulate large-scale semi-structured data, like JSON data
  - Parallelism: uses the Hadoop MapReduce framework to process JSON data in parallel
  - Extensibility: JAQL UDFs, Java functions, JAQL modules
- JAQL provides a hook into any BigInsights analysis

HDFS_READ Example

HDFS_READ

Read a file from HDFS, present the file as a DB2 table for use in SQL

```
SET RESULT_FILE = JAQL_SUBMIT(. . .

SELECT BIRESULT.CNT FROM
TABLE(HDFS_READ(RESULT_FILE, '')

AS BIRESULT(CNT INTEGER);
```

URL of the CSV file to be read

options

Definition of the “table”, how to present the results to SQL

Integrated Query Example

```
INSERT INTO BI_TABLE (CNT)
(SELECT CNT FROM TABLE
(HDFS_READ
 (JAQL_SUBMIT
  ('syslog = lines("hdfs:///idz1470/syslog3sec.txt");

    [read(syslog)->count()]->
     filter(strPos($,"$HASP373")>=0)->
     write(del(location="hdfs:///idz1470/iod00s/lab3e2.csv"));',

   'http://bootcamp55.democentral.ibm.com:14000/webhdfs/v1/idz1470/
    iod00s/lab3e1.csv?op=OPEN',

   'http://bootcamp55.democentral.ibm.com:8080',
   ''
  ),
  ''
 )
)
AS BIGINSIGHTS(CNT INTEGER));
```

JAQL_SUBMIT can be embedded in HDFS_READ for a synchronous execute/read workflow

New - BigInsights Try and Buy Program

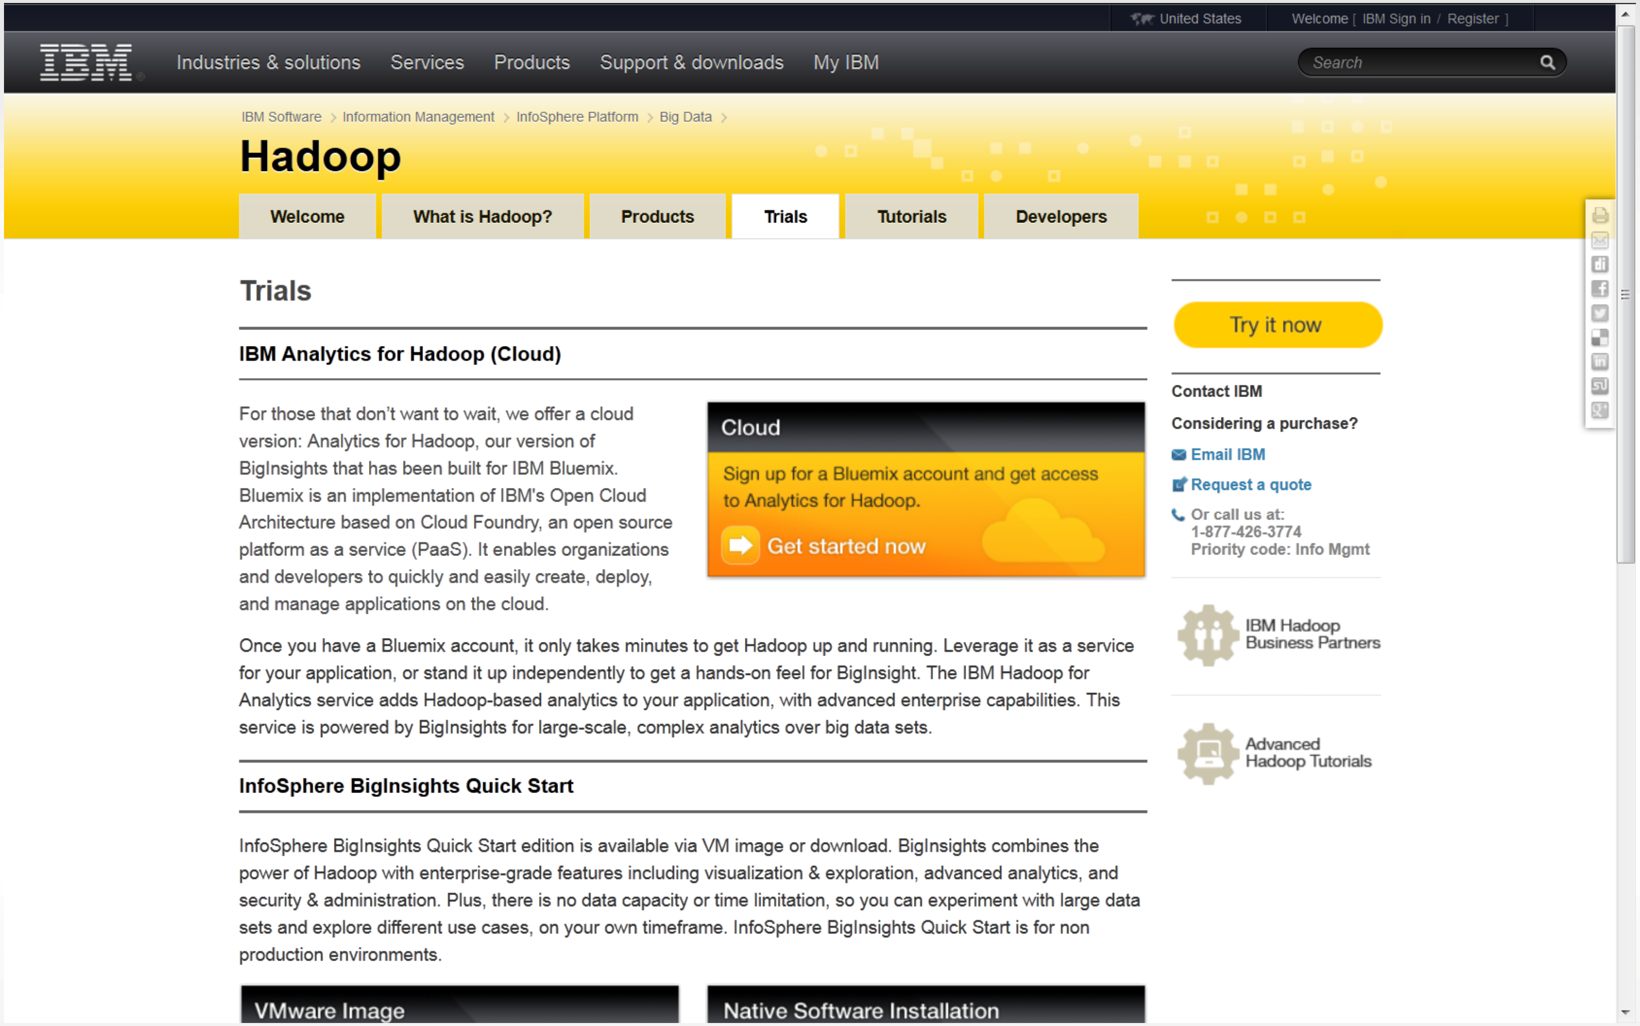

Hadoop

Trials

IBM Analytics for Hadoop (Cloud)

For those that don't want to wait, we offer a cloud version: Analytics for Hadoop, our version of BigInsights that has been built for IBM Bluemix. Bluemix is an implementation of IBM's Open Cloud Architecture based on Cloud Foundry, an open source platform as a service (PaaS). It enables organizations and developers to quickly and easily create, deploy, and manage applications on the cloud.

Cloud

Sign up for a Bluemix account and get access to Analytics for Hadoop.

Once you have a Bluemix account, it only takes minutes to get Hadoop up and running. Leverage it as a service for your application, or stand it up independently to get a hands-on feel for BigInsight. The IBM Hadoop for Analytics service adds Hadoop-based analytics to your application, with advanced enterprise capabilities. This service is powered by BigInsights for large-scale, complex analytics over big data sets.

InfoSphere BigInsights Quick Start

InfoSphere BigInsights Quick Start edition is available via VM image or download. BigInsights combines the power of Hadoop with enterprise-grade features including visualization & exploration, advanced analytics, and security & administration. Plus, there is no data capacity or time limitation, so you can experiment with large data sets and explore different use cases, on your own timeframe. InfoSphere BigInsights Quick Start is for non production environments.

Contact IBM

Considering a purchase?

Email IBM

Request a quote

Or call us at:
1-877-426-3774
Priority code: Info Mgmt

IBM Hadoop Business Partners

Advanced Hadoop Tutorials

Agile Development and DevOps

Software delivery is at the heart of today's top technology trends

Big Data
Insights on new products by more efficiently interpreting massive quantities of data

Cloud
Demand for apps requires fast, scalable environments for dev and test, as well as production

Social Business
Broader set of stakeholders collaborates to deliver continuous innovation and value

Instrumented Products
Industry requirements demand faster response to regulations and standards, with traceability and quality

Mobile
Modern workforce expects constantly updated software to connect to enterprise systems

Intelligent/ Connected Systems
Software component in smart products driving increased value and differentiation

Software delivery

…and yet, a lack of continuous delivery impacts the entire enterprise in the new reality of Systems of Interaction

Line-of-business
Takes too long to introduce or make changes to mobile apps and services

Operations
Rapid app releases impact system stability and compliance

>70%

of budgets devoted to maintenance and operations

4-6 weeks

to deliver even minor application changes to customers

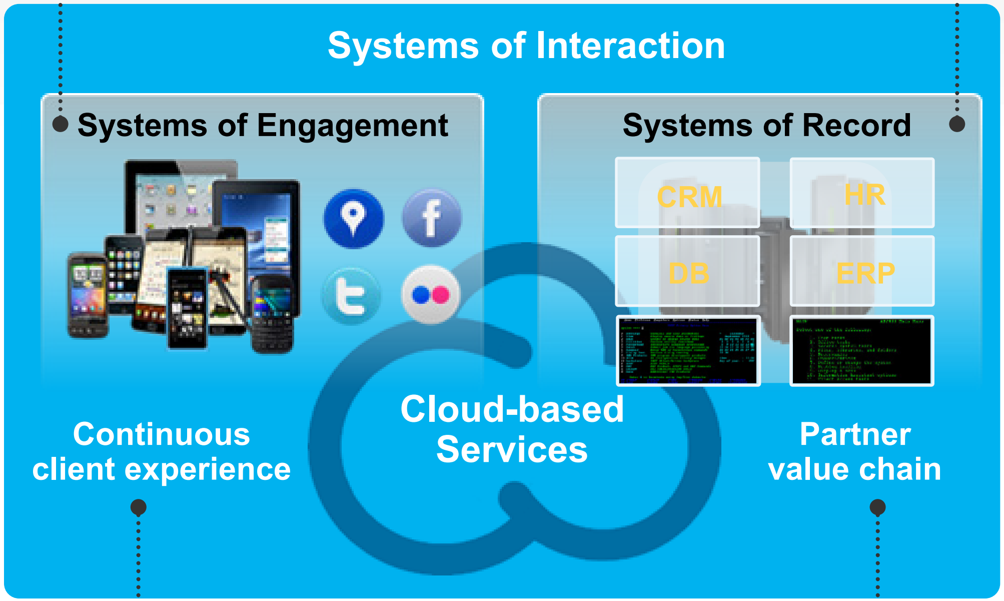

>45%

of customers experience production delays

>50%

of outsourced projects fail to meet objectives

Development/Test
Speed mismatch between faster moving front office and slower moving back office systems, delaying time to obtain feedback

Suppliers
Delivery in the context of agile

Enabling Agile programming with the DB2 JSON API

NoSQL Overview

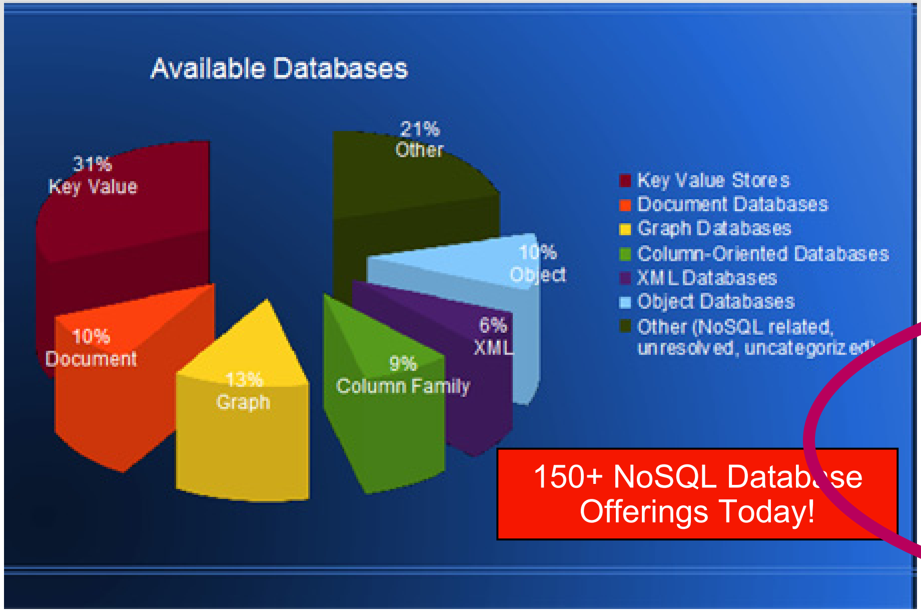

150+ NoSQL Database Offerings Today!

Motivation

- Many apps need fewer database features (simplicity)
- Need rapid application evolution/deployment, with minimal interaction with DBA
- Some apps need extremely high scale (e.g. Twitter)
- Need for a low-latency, low-overhead API to access data
- Increasing use of distributed analytics

Dominant Flavors

- Key Value Stores
  - Hash table of keys, where the data part of key-value is in a binary object.
  - Examples pure key-value stores: MemcacheD, REDIS, WebSphere eXtreme Scale.
- Document Stores
  - Stores documents made up of tagged elements, which have keys and document-like objects.
  - Examples: MongoDB, couchDB.
- Column Family
  - Each storage block contains data from only one column/column set.
  - Examples: Hbase, Cassandra.
- Graph Store
  - Key-values are related through graph structure.
  - Common Model: *RDF.*
  - Examples: *Jena, Sesame*

JSON is the Language of the Web

- **JavaScript Object Notation**
  - Lightweight data interchange format
  - Specified in IETF RFC 4627
  - http://www.JSON.org
- Designed to be minimal, portable, textual and a subset of JavaScript
  - Only 6 kinds of values!
  - Easy to implement and easy to use
- Text format, so readable by humans and machines
- Language independent, most languages have features that map easily to JSON
- Used to exchange data between programs written in all modern programming languages

```
{
    "firstName“         : "John",
    "lastName"          : "Smith",
    "age"               : 25,
    “active”            : true,
    “freqflyer_num      : null,

    "address" :
    {
      "streetAddress“: "21 2nd Street",
      "city"          : "New York",
      "state"         : "NY",
      "postalCode"    : "10021"
    },

    "phoneNumber“     :
    [
      {
        "type"     : "home",
        "number“: "212 555-1234"
      },
      {
        "type"     : “mobile",
        "number“: "646 555-4567"
      }
    ]
 }
```

New Era Application Requirements

- **Store data from web/mobile apps in it's native form**
  - New web and mobile applications use JSON for storing and exchanging information
    - Very lightweight – write more efficient apps
    - Preferred data format for mobile application back-ends

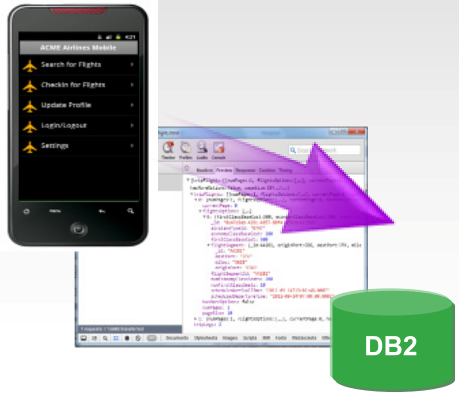

- **Move from development to production in no time!**
  - Ability to create and deploy flexible JSON schema
  - Gives power to application developers by reducing dependency on IT; no need to pre-determine schemas and create/modify tables
  - Ideal for agile, rapid development and continuous integration

Common JSON Open Source Data Store Attributes

- Logging often turned off to improve performance

- By default, no return code on insert (a.k.a. “fire and forget”)
  - App must verify update was performed

- Data is sharded for scalability
  - Shards are replicated asynchronously for availability
  - Provides “Eventual Consistency” - Queries to replica nodes can return back-level data

- No concept of commit or rollback
  - Each JSON update is independent
  - Applications have to implement compensation logic to update multiple documents with ACID properties

- JSON documents are stored in collections
  - But no “join” across collections

- No document-level locking
  - App must manage a “revision” tag to detect document update conflicts

DB2 for z/OS Enterprise-class JSON Database
Agility with DB2 Qualities of Service

- **Combine data from “systems of engagement” with core enterprise data**
  - Simplicity and agility of JSON + enterprise strengths of DB2
  - Simplify data access when a single application’s data has different access characteristics

- **Maintains JSON simplicity and agility**
  - Interoperate seamlessly with modern applications
  - Flexible schemas allow rapid delivery of applications

- **Leverages DB2 Qualities of Services**
  - Security
  - Management, operations
  - High availability

- **Delivers the best of both worlds**

DB2 for z/OS JSON Document Store

Announced 10/1, Generally Available 12/6

- **Java Driver for JSON API**
  - Java Driver supporting JSON API
  - Transactions
  - Parametric SQL statements (Delete, Select)
  - Temporal tables (DB2 LUW)
- **JSON Wire Listener**
  - Leverage community drivers
- **CLP-Like Command Shell**
  - Ad-hoc updates / queries
  - Administration commands
- **DB2 Data Server**
  - JSON documents stored as BSON (Binary JSON) in BLOBs within DB2
  - Scalar function and UDF extensions
  - Enhanced indexing on expression capabilities allows indexing of JSON members

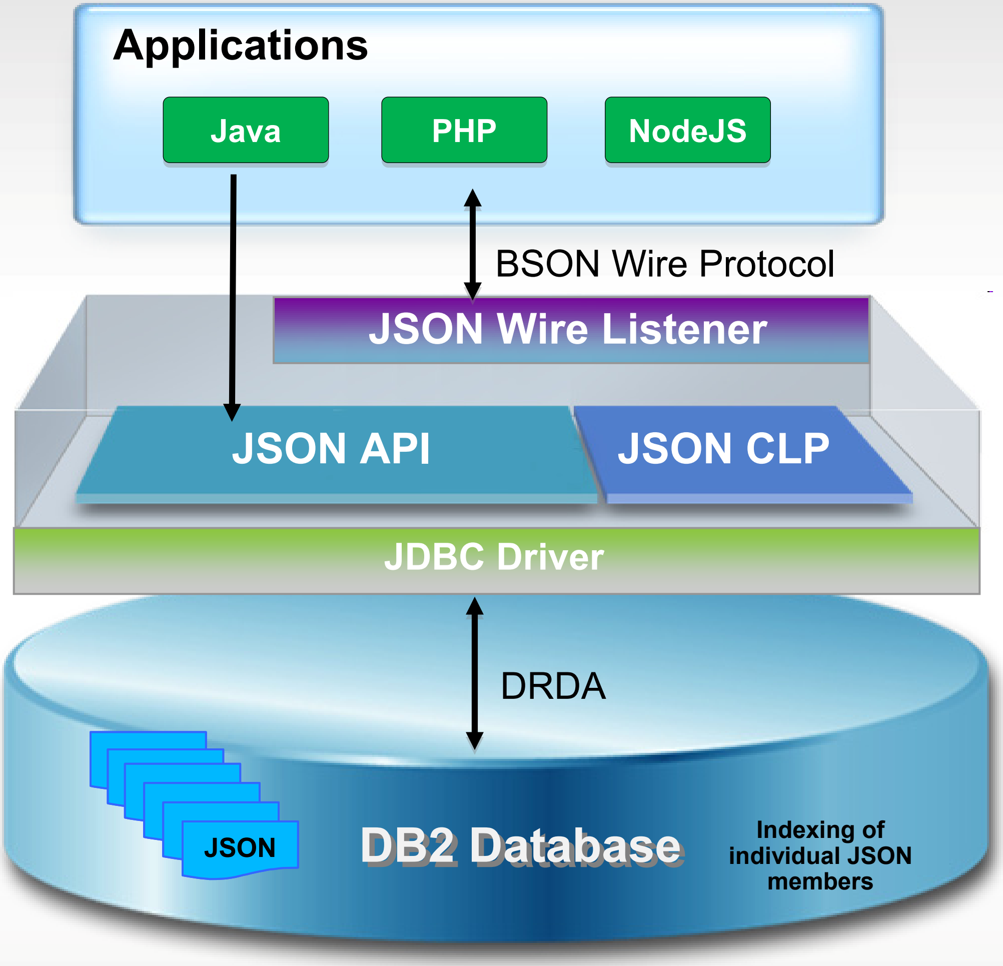

Java Driver for JSON API

Translates API calls to SQL + function invocations

- Implemented MongoDB API – simple, easy to use
- Supports transactions
- Batches insertions
- Fire-forget inserts (fast)
- Indexing

Insert a record, a blog post by Joe:

```
db.posts.insert({author:"Joe", date:"2012-04-20", post:"..."})
```

Find all posts by Joe:

```
db.posts.find({author:"Joe"})
```

Delete all posts of Joe:

```
db.posts.remove({author:"Joe"})
```

JSON API Core Architecture

***Java Driver** that translates API calls to SQL + function invocations*

- Supports Transactions
- Batches insertions
- Fire-forget inserts (fast)
- Indexing
- Time travel query (DB2 LUW)
- Smart Query re-write
- Inline LOBS allow buffering and compression

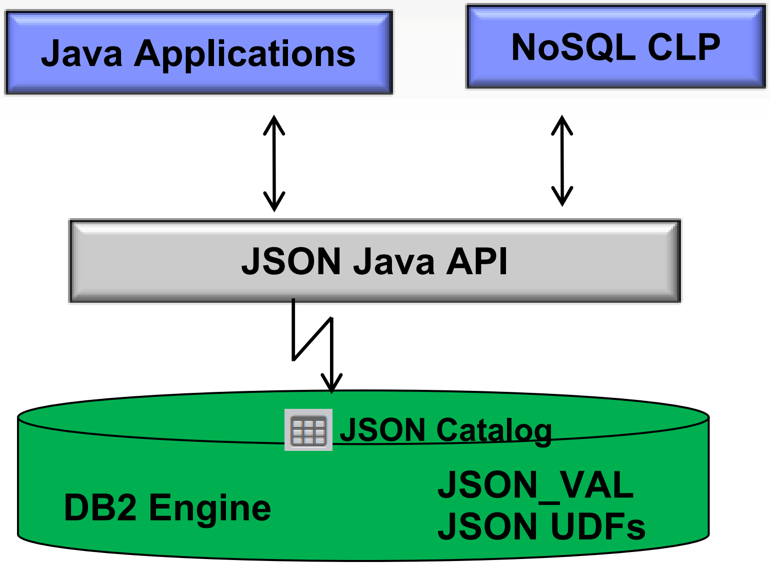

- Java command shell
  - Administration
  - Ad-hoc Query and Update

JSON API Calls - Under the Covers

1) Create a customer collection / table.

```
db.createCollection(“customers")
```

```
CREATE TABLE customers (ID VARBIN(12),
                        DATA BLOB(16MB)
                        INLINE(25K))
* Also registers collection in SYSJSON_INDEX
```

2) Insert all your customers as JSON documents. For example, one insert might contain this document:

{ name:"Joe", age:25, phone:["555-666-7777", "444-789-1234"],
address: { street:"ABC st",
zipcode:"95141“ } }

```
db.customers.insert({name:”Joe”…)
```

```
INSERT INTO customers (DATA)
       VALUES ( <binary JSON> )
```

3) Look for customers in zipcode 95141.

```
db.customers.find(
     {“address.zipcode”:’95141’})
```

```
SELECT DATA FROM customers
WHERE JSON_VAL
       (DATA,'address.zipcode','s:5')
      ='95141'
```

4) Improve performance by creating index on zipcode.

```
db.customers.ensureIndex
   ({“address.zipcode"}});
```

```
CREATE INDEX idx1
ON customers
 (JSON_VAL(DATA,'address.zipcode','s:5')

* Also registers index in SYSJSON_INDEX
```

DB2 JSON Store – NoSQL JSON Wire Listener

•Immediate reach to more applications and developers

- Leverage community drivers
- Presence in “New style apps”
- Built on JSON Java API
- Area for IBM collaboration with MongoDB, open source community

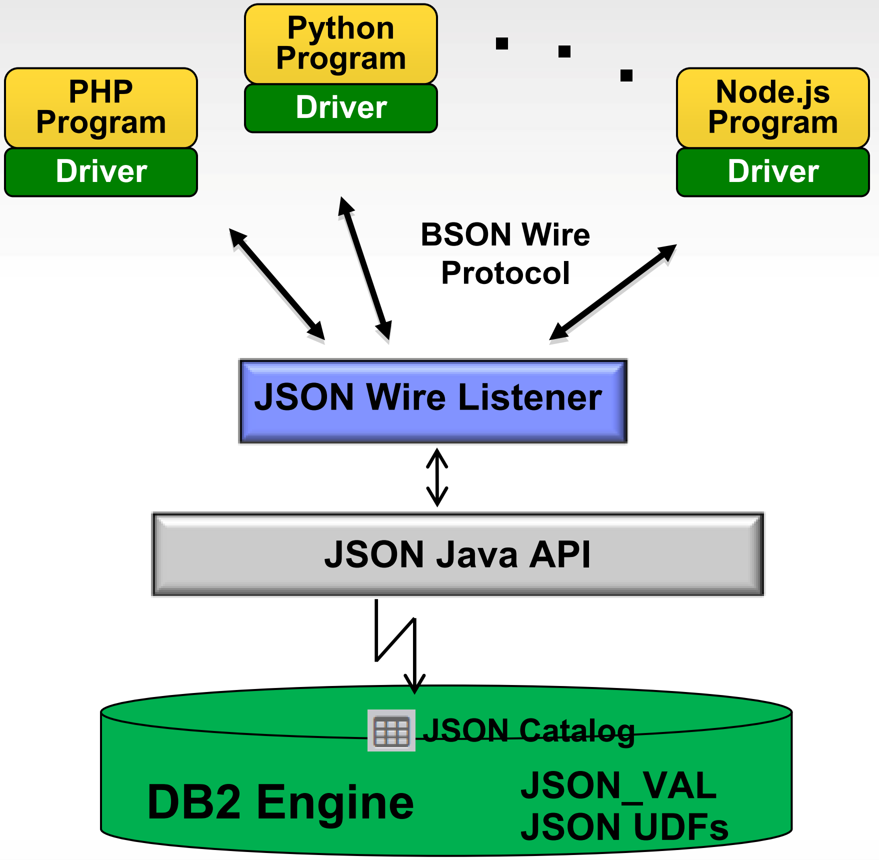

What data store format makes sense for your application?

- **Consider DB2 JSON when:**
  - Application and schema subject to frequent changes
  - Prototyping, early stages of application development
  - De-normalized data has advantages
    - Entity / document is in the form you want to save
    - Read efficiency – return in one fetch without sorting, grouping or ORM mapping
  - Systems of Engagement
    - Less stringent "CAP" requirements in favor of insert speed
    - Social media

- **Relational still best suited when these are critical:**
  - Data normalization to
    - Eliminate redundancy
    - Ensure master data consistency
  - Database-enforced constraints

Future Trends – A Look Ahead

- DB2 JSON API
  - Fill in Functional Gaps with MongoDB Capabilities
  - Performance, Performance, Performance
  - Extend DB2 Capabilities to the Wire Listener

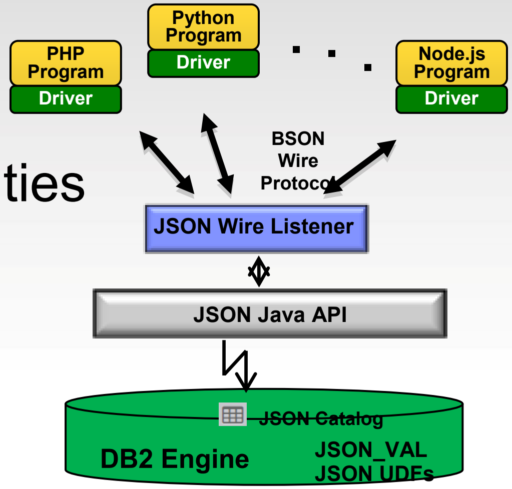

- JSON in SQL
  - *** Possible Future Direction ***
  - The DB2 JSON API uses SQL under the covers, as shown
    - To move that to a supported feature?
      - JSON / BSON conversion
      - Improved friendliness of JSON built-in function and UDFs
      - Reconciliation with SQL Standard

JSON in SQL – First Steps

Extend JSON API Building blocks for external use

- New functions released in DB2 11 only
  - JSON2BSON - convert JSON string into BSON format
  - BSON2JSON - convert BSON LOB into JSON string
  - JSON_VAL - retrieve specific value from inside a BSON object

- Examples:

```
CREATE TABLE FLIGHTLOG (FLIGHTNUM CHAR(4), FLIGHTDETAILS BLOB)

INSERT INTO FLIGHTLOB (‘1234’, BLOB(SYSTOOLS.JSON2BSON(‘{departure: {date: new Date(…’})))

SELECT SYSTOOLS.BSON2JSON(FLIGHTDETAILS) FROM FLIGHTLOG WHERE FLIGHTNUM = 1234

SELECT SYSTOOLS.BSON2JSON(FLIGHTDETAILS) FROM FLIGHTLOG WHERE
      JSON_VAL(FLIGHTDETAILS, ’departure,date’, ‘d’) = ’12/21/2014/)
```

Rational Developer For System z (RDz)

DevOps takes a closed-loop approach to application delivery

Extending Lean and Agile practices across the entire software value chain

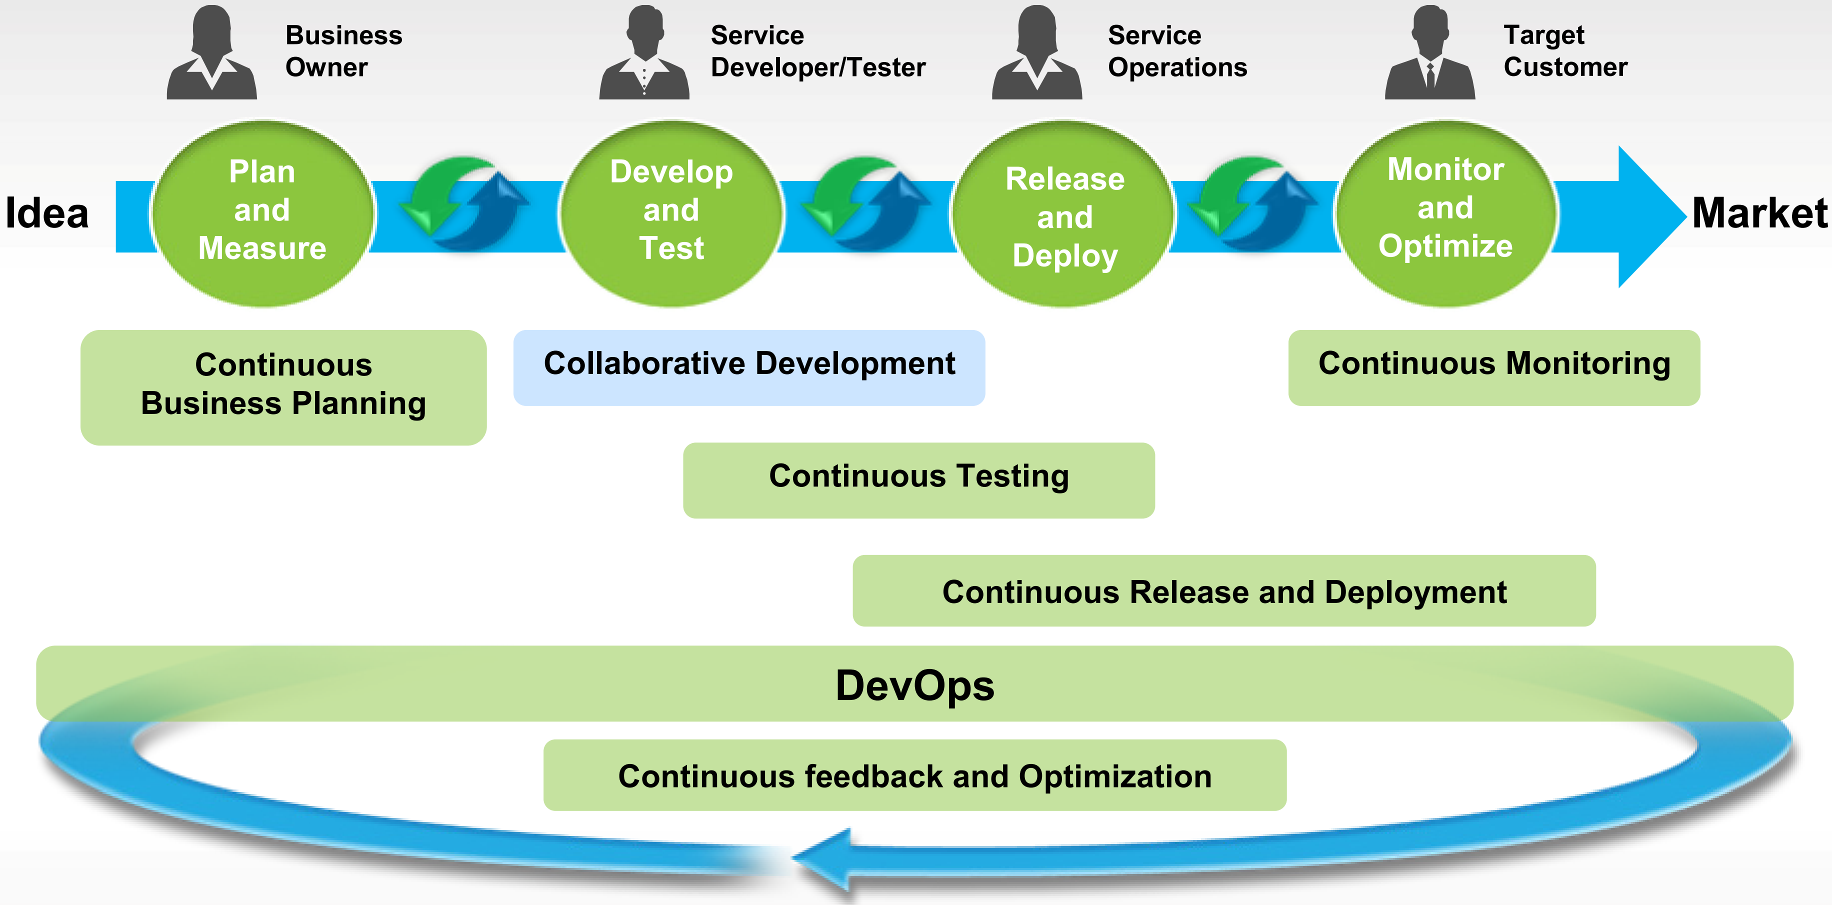

balance speed, cost, quality and risk

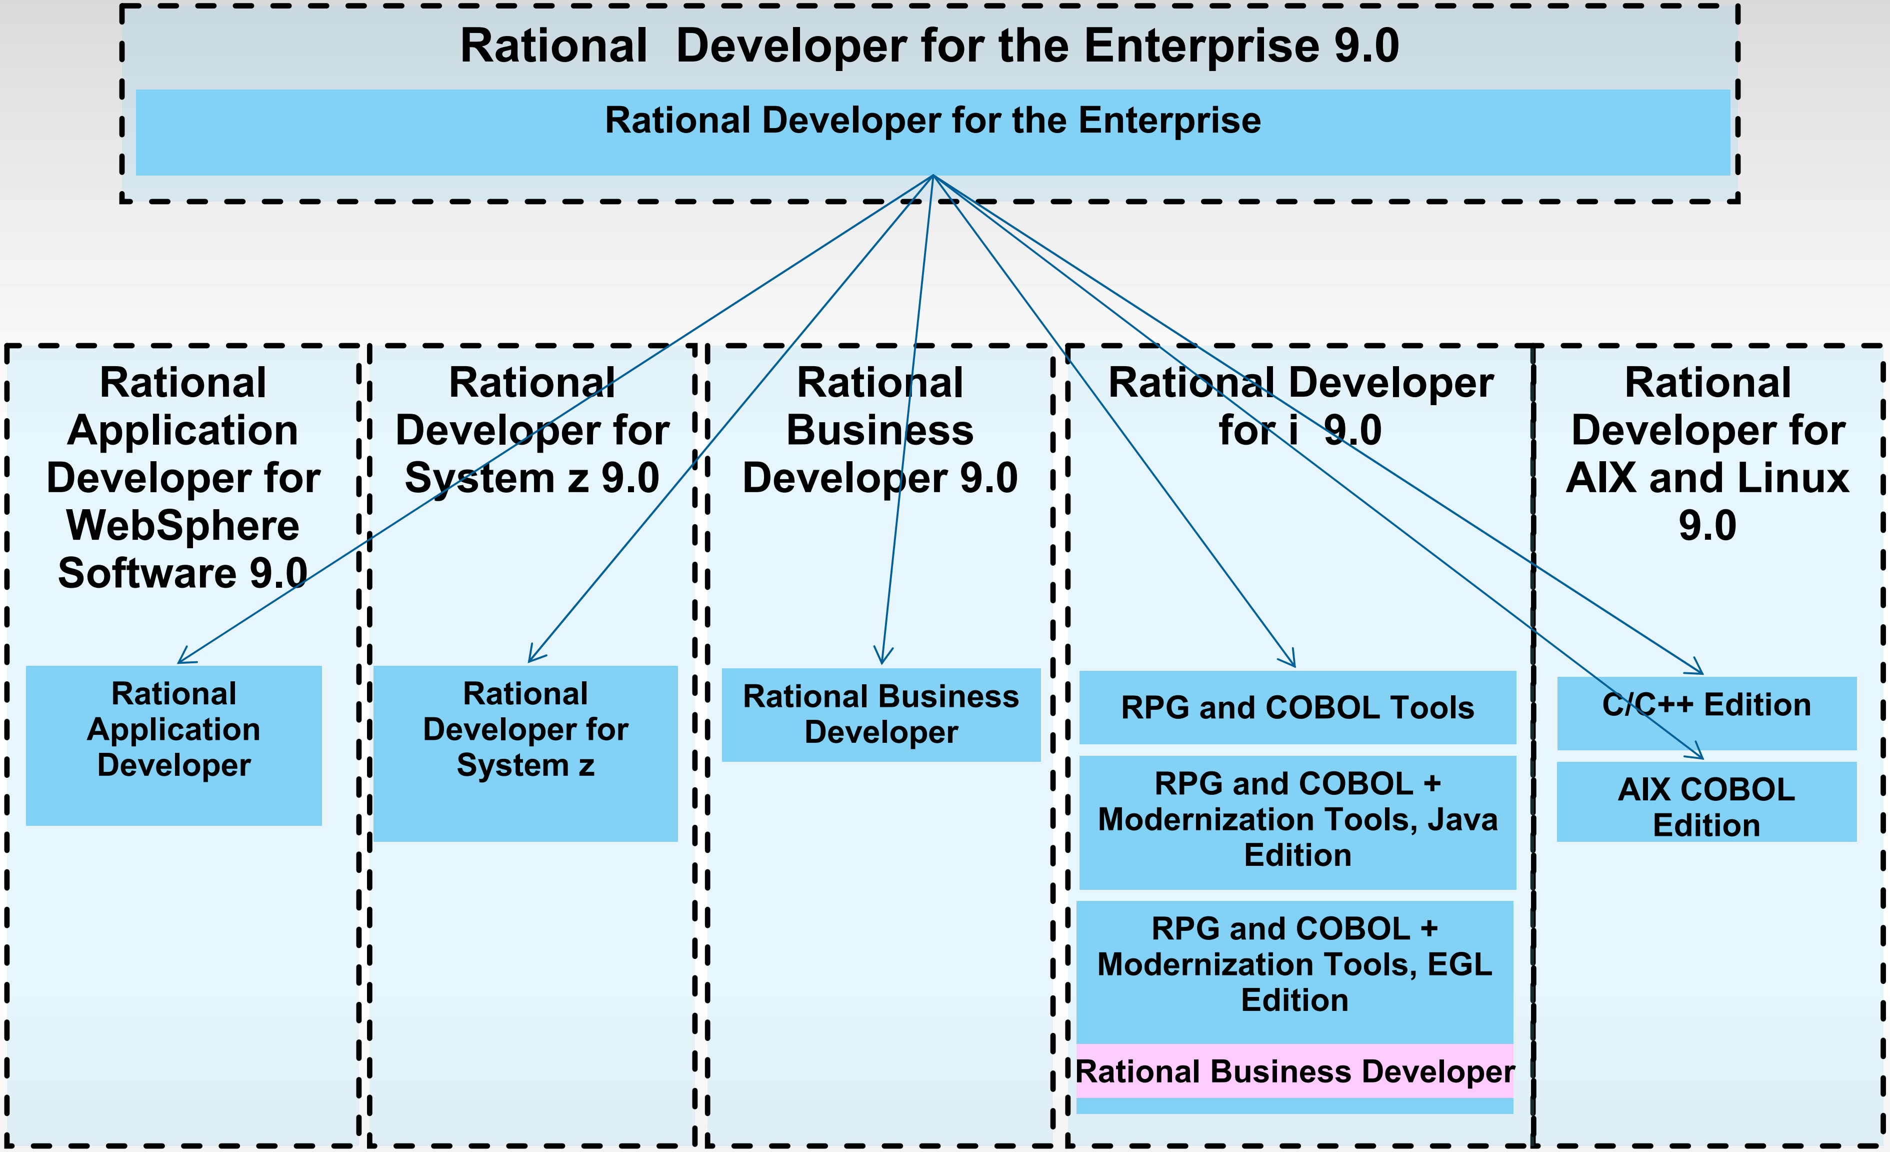

Rational Developer for the Enterprise 9.0

Rational Developer for the Enterprise

Rational Application Developer for WebSphere Software 9.0

- Rational Application Developer

Rational Developer for System z 9.0

- Rational Developer for System z

Rational Business Developer 9.0

- Rational Business Developer

Rational Developer for i 9.0

- RPG and COBOL Tools
- RPG and COBOL + Modernization Tools, Java Edition
- RPG and COBOL + Modernization Tools, EGL Edition
- Rational Business Developer

Rational Developer for AIX and Linux 9.0

- C/C++ Edition
- AIX COBOL Edition

What are their foundations?

eclipse
JUNO
(c) Copyright Eclipse contributors and others, 2000, 2012. All rights reserved. Eclipse is a trademark of the Eclipse Foundation, Inc. Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

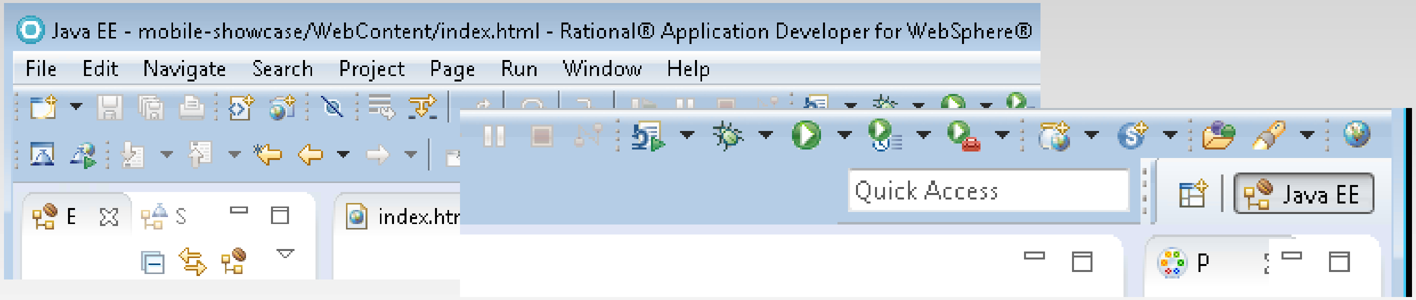
Java EE - mobile-showcase/WebContent/index.html - Rational® Application Developer for WebSphere®
File
Edit
Navigate
Search
Project
Page
Run
Window
Help
E
S
index.htr
Quick Access
Java EE
P

- **Eclipse 4.2.2**
  - Modern visual style, improved toolbar presentation
  - Highly extensible & customizable user interface
  - Flexible layouts with detachable editors
  - Global search capability, improved navigation, shortcuts, validations, tracing
  - Many more enhancements in Platform, JDT, Equinox & PDE

- **New Desktop Operating Systems support**

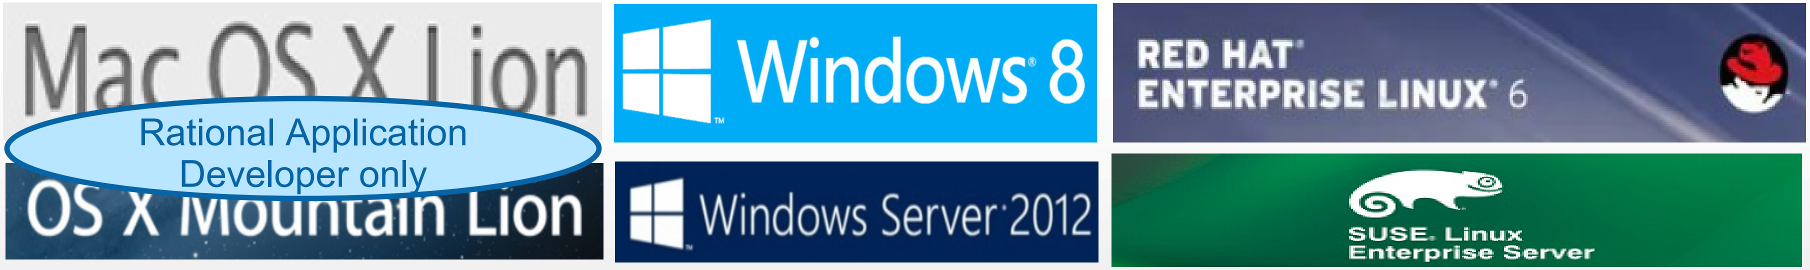

Rational Application Developer only

IBM Rational Developer for System z

An integrated development environment enabling enterprise modernization:

- ✓ An Eclipse-based IDE designed to improve the productivity of mainframe developers regardless of their technical background
- ✓ Tooling to create, modernize and maintain COBOL, PL/I, C/C++, Java, or Assembler transactional applications for CICS, IMS, and IBM DB2
- ✓ Support for SOA and Web services creation
- ✓ Integration with the Rational product suite in support of mobile development and DevOps
- ✓ Cross-platform support (z/OS, AIX, Linux on Power, Linux x86 and IBM i) provided by RD for the Enterprise

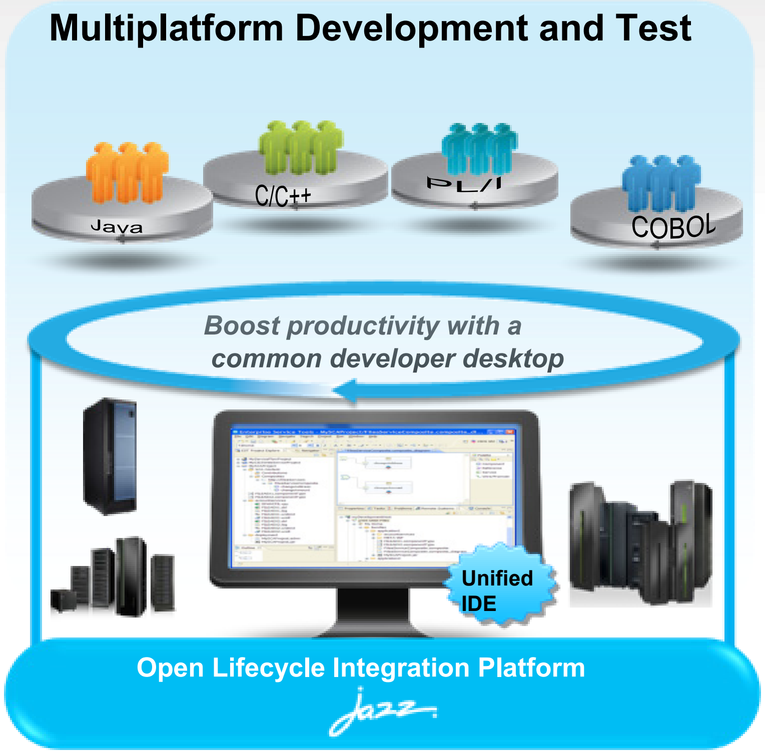
Multiplatform Development and Test
Java
C/C++
PL/I
COBOL
Boost productivity with a common developer desktop
Unified IDE
Open Lifecycle Integration Platform
Jazz

RDz V9.0.1 Features

- ✓ Integrated Debugger and Code Coverage
- ✓ RDz/RTC integration – User build improvements
- ✓ JES interface improvements
- ✓ Remote z/OS Search enhancements
- ✓ Migrated data set enhancements
- ✓ Editor updates (JCL, COBOL, PLI, System z Lpex, System z Data editor)
- ✓ Enterprise Service Tools enhancements
- ✓ Code review enhancements including New PLI Code rules

New SQL and DB2 Application Development Features

Transparent Archive Data

Access “warm” and “cold” data in a single query

Five easy steps:

1. Create a base table:

```
CREATE TABLE POLICY_ACTIVE (     POLICY_ID CHAR(10) NOT NULL,
                                 COVERAGE INTEGER NOT NULL);
```

2. Create an archive table:

```
CREATE TABLE POLICY_ARCHIVE (  POLICY_ID CHAR(10) NOT NULL,
                               COVERAGE INTEGER NOT NULL);
```

3. Enable transparent archiving:

```
ALTER TABLE POLICY_ACTIVE ENABLE ARCHIVE USE POLICY_ARCHIVE;
```

4. Before deleting rows, Tell DB2 to transparently archive base data – set built-in global variable, SYSIBMADM.MOVE_TO_ARCHIVE:

```
SET SYSIBMADM.MOVE_TO_ARCHIVE = ‘Y’;      /* data will be archived */
DELETE FROM POLICY_ACTIVE WHERE...;
```

5. To retrieve both current and archived data, use built-in global variable SYSIBMADM. GET_ARCHIVE (or the ARCHIVESENSITIVE BIND option)

```
SET SYSIBMADM.GET_ARCHIVE = ‘Y’;
SELECT COUNT(*) ... FROM POLICY_ACTIVE
```

Global Variables

- Simple, yet powerful feature for cross-SQL communication
  - No need to retrieve state data back to application, then send on next SQL
- Example:

```
CREATE VARIABLE SESSION_START_TS TIMESTAMP
                       DEFAULT CURRENT TIMESTAMP;
```

- User-defined name, type, default value
- Scope is similar to special registers – connection scope, not shared across sessions.

Array Data Type

- Convenient storage an lookup structure
- Allows more efficient invocation of stored procedures with variable tuples
- Permitted as:
  - An SQL variable
  - A parameter or **RETURNS** data-type of an SQL scalar function
  - A parameter of a **native** SQL procedure
  - The target data type for a **CAST** specification
- Examples:
  - Regular Array:

    ```
    CREATE TYPE PRODUCT_NUMS AS CHAR(8) ARRAY[30];
    ```

  - Associative Array:

    ```
    CREATE TYPE PRODUCT_NAMES AS VARCHAR(40)  ARRAY[CHAR(8)];
    ```

Additional Miscellaneous Enhancements

- Grouping Sets, including ROLLUP and CUBE
- Temporal Data additions
- Autonomous SQL PL Procedures
- In-database real time scoring for SPSS
- Suppressing Null Indexes
- XML Enhancements
  - XQuery support
  - XML performance enhancements in DB2 10 and DB2 11
  - XQuery FLWOR expressions performance enhancements
  - XMLTABLE performance enhancements in DB2 11

DB2 Connect and Language Interface Update

Dynamic Language Interfaces to DB2

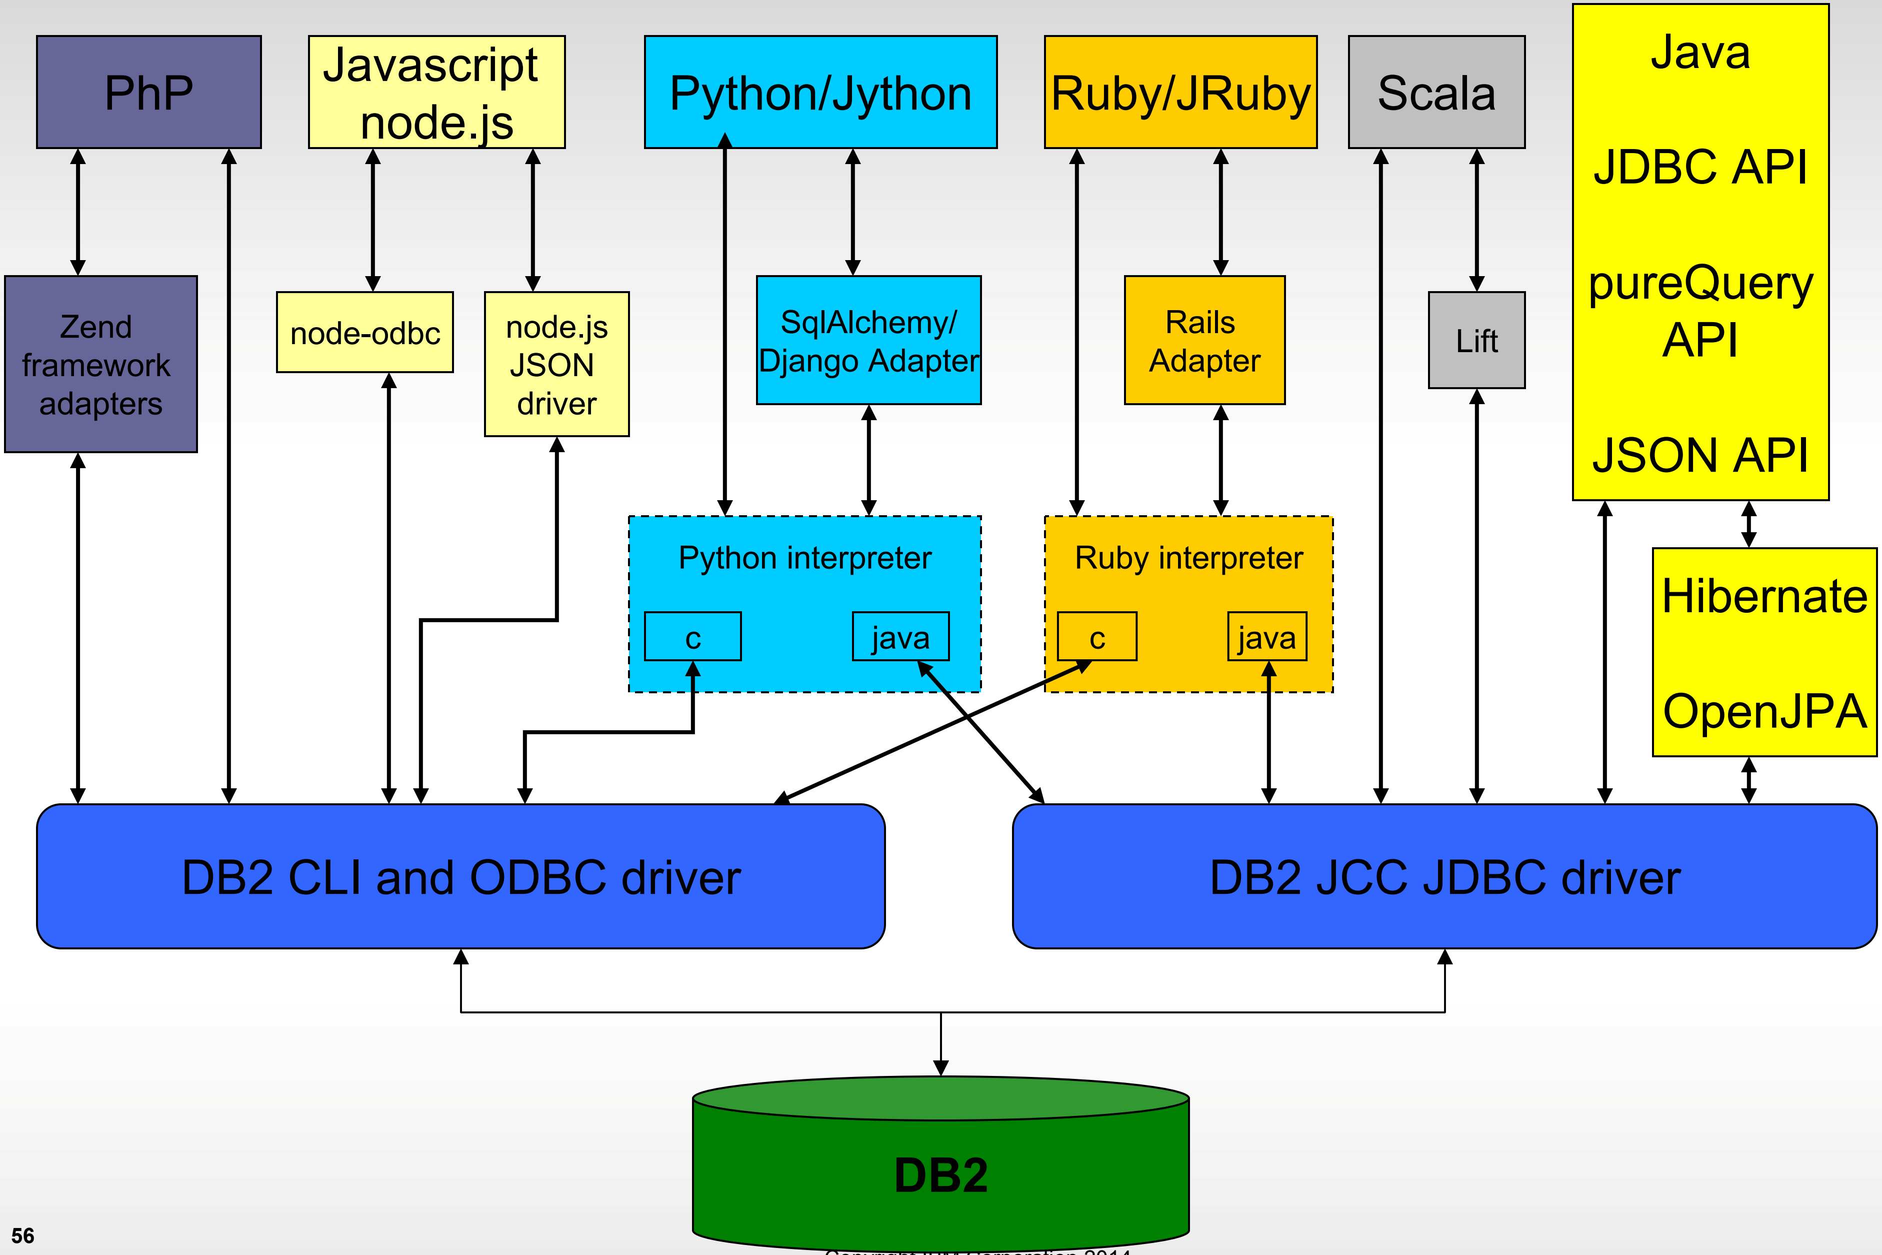

Recent Connect Enhancements Summary

- DB2 Connect
  - Fixpack 2 – DB2 z/OS 11, new SQL support:
    - CALL with array
    - Global Session Variables
    - Larger Client Info Fields
    - more…
  - Fixpack 3 – JSON API support for DB2 z/OS
  - Fixpack4 – License serviceability, Seamless Failover

CLI for z/OS – Recent Enhancements

- **Fixpack 4**
  - **Enhancements to the seamless failover feature**
    - SQL_ATTR_REPORT_SEAMLESSFAILOVER_WARNING connection attribute shows whether to return a warning message when a connection fails over seamlessly.
    - detectReadonlyTxn keyword to modify the seamless failover feature
  - **Control array input chaining for multi-row INSERT**
    - SQL_ATTR_COLUMNWISE_MRI attribute and ColumnwiseMRI keyword for the MERGE statement enables array input chaining.
  - **General enhancements to the CLI driver**
    - GRANT and GRANT_ROLE bind options now allowed on SQLCreatePkg() function and db2cli commands
    - SQL_AUTOCOMMIT_DEFERRED value for the SQL_ATTR_AUTOCOMMIT connection attribute.
    - ClientWrkStnName or ClientWorkstationName keyword now allow the 'NODEFAULT' value
      - prevent a domain name server (DNS) lookup when the client workstation value is not required.
    - Obtain a list of all configuration keywords that are set on the server
    - Control retry of MERGE stmt up to 6 times
    - Specify a read-only connection

- **Fixpack 2:**
  - **Array Parameters on CALL statements**
  - **Support implicit COMMIT operation for stored procedure calls**
    - avoid extra network trip
  - Set Client Correlation Token information
  - Support Global Variables

JDBC enhancements for z/OS

- **DB2 Cancun Release 10.5.0.4: Alternate group support**
  - Alternate group support allows the IBM Data Server Driver for JDBC and SQLJ to move an application workload to a DB2 for Linux, UNIX, and Windows alternative DB2 pureScale instance when the primary group is unavailable.

- **DB2 Cancun Release 10.5.0.4: Methods for converting ResultSet objects to JSON format**
  - You can query a relational table, and then convert the ResultSet that you retrieve to a com.ibm.db2.jcc.json.DB2JSONResultSet object, which holds the data in JSON format. After you create the DB2JSONResultSet object, you can use methods on that object to retrieve the JSON documents or JSON snippets.

- **DB2 Cancun Release 10.5.0.4: Improved implementation of ResultSet.getXXX for DECIMAL columns**
  - The IBM Data Server Driver for JDBC and SQLJ implementation of ResultSet.getXXX methods that return numeric values has been improved. As a result, you might experience better performance when you retrieve DECIMAL column data from a table.

- **DB2 Cancun Release 10.5.0.4: Enhancements that provide more DB2 Connect client license information**
  - If you use a DB2 Connect client license for connections to DB2 for z/OS or DB2 for i, you can use IBM Data Server Driver for JDBC and SQLJ enhancements to obtain the following information about the client license:
    - License expiration date
    - License type
    - DB2 Connect version that is valid with the client license

Additional Resources for the DB2 JSON API

Read 4-part DB2 JSON Tech Article Series

- Introduction to DB2 JSON
  ibm.co/15ImEke
- Command line processor
  ibm.co/GYfi3e
- Writing apps with Java API
  ibm.co/19RWv5Y
- JSON Wire Listener
  ibm.co/16aLEmF

Announcement Details (z/OS)

- DB2 for z/OS Accessories Suite
  http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ca&infotype=an&supplier=897&letternum=ENUS213-395

Additional Resources on the BigInsights Connector

- BigInsights (including JAQL documentation)
  - http://www-01.ibm.com/software/data/infosphere/biginsights/