

# DB2 10.5 for zLinux

## New Features



**Stefan Hummel**  
Senior IT Specialist

---

## Legal Notice

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 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

- **The History**
- **DB2 10.5 for LUW – What's new, what's changed  
(Focus on zLinux)**
- **Best Practices**

# DB2: Proven to Tens of Thousands of Clients



The Ultimate Driving Machine®





# History of DB2 for zLinux

- In 2000, **DB2 7.1** was one of the first software programs certified for use on Linux for System Z
- The IBM DB2 and Linux for System Z development teams work together on the DB2 on Linux for System Z stack
- Quality, reliability, and performance **proven by hundreds of satisfied DB2 on Linux for System Z customers**
- In 2013 IBM delivered **DB2 10.5** for Linux on System Z

**More than**

**10 years of Enterprise Linux® on System z®**

A simple idea that changed the world

**running**  
**DB2**



2000 2001 2002 2003 2004 2005 2006 2007 2008 2009 2010

## Choose from 2 Linux for System Z Distributions

- IBM is committed to open standards Linux
- DB2 is **optimized to run on SuSE and Red Hat Linux** distributions



- Favorite levels
  - SLES 11 SP2
  - RHEL 6.3

# Recent DB2 Evolution

	<b>DB2 9.1</b> <i>(mid '06)</i>	<b>DB2 9.5</b> <i>(late '07)</i>	<b>DB2 9.7</b> <i>(mid '09)</i>
<b>Total Cost of Ownership</b>	Auto Storage Mgt Adaptive Memory Mgt	Auto Compression Fully Automatic Memory Mgt Automatic Statistics Mgt Advances Single System View Cluster DB Mgt	360° Monitoring In memory SQL based Reorg Avoidance
<b>Continuous Availability</b>	Advanced Recovery Restartable Rebuild DB via Tbsp images	Integrated HA and DR Integrated Flash Copy Hardware Memory Protection	HADR Advances Active Standby Trap Resilience General Online Schema Evolution
<b>Performance &amp; Scale</b>	Deep Compression Range Partitioning Limit Removal	Integrated Workload Mgt New Multi-threaded Architecture Optimized Bulk Deletion	Deeper Deep Compression Scan Sharing WLM Advances
<b>Security &amp; Privacy</b>	Label Based Access Control	Database Roles Identity Assertion Advanced Auditing	Separation of Duties
<b>Application Enablement &amp; XML</b>	pure XML	Optimistic Locking Decimal Floating Point Pure XML Advances	XML extended to DPF, MDC, RP "Currently Committed" Isolation Numerous SQL Features

# Proven Compelling DB2 9.7 Themes

## ■ Resource Optimization

- Best performance with most efficient utilization of available resources
  - Index compression
  - Inline LOBs

## ■ Operational Flexibility

- Allow for continuous and flexible change management
  - Schema evolution
  - Local indexes for range partitioned tables
  - Tablespace management

## ■ Service Level Confidence

- Expand your critical workloads confidently and cost effectively
  - WLM (Workload Management)

## ■ XML Insight

- Harness the business value of XML


## ■ Break Free with DB2

- Oracle compatibility features



# Up to 10x Savings with Adaptive Compression

Lower Storage Costs; Lower Administration Costs



- Adaptively apply both table-level compression and page-level compression
  - Unique capabilities in the industry
- Table re-orgs not required to maintain high compression
- Compress archive logs

# DB2 10 and InfoSphere Warehouse 10

## *DB2 is at the core of both*



DB2 10 is the  
underlying  
database!

### DB2 10

DB2 database software offers industry **leading performance and reliability** on a choice of platform from Linux, Unix and Windows

### InfoSphere Warehouse 10

A complete real-time data warehousing platform that delivers **superior scalability and availability**, design, build, and management tooling, and business analytics.

## DB2 10.5 for zLinux – What's new ?

- **Oracle Application Compatibility (cont.)**
  - Continue to reduce the cost and risk associated with migrating Oracle applications to DB2
  
- **Mobile**
  - Rich capabilities to support mobile devices
  
- **NoSQL**
  - Continue to support the next generation of applications
  
- **Enhanced Tooling**
  - Reducing the total cost of ownership with DB2 and making the adoption, management, monitoring, and maintenance very simple

# PL/SQL Compatibility Above 98%

## Easily Move from Oracle Database

### Reliance Life Insurance

*“The total cost of ownership with DB2 running on IBM systems is almost half the cost of Oracle Database on Sun systems.”*

### Banco de Crédito del Peru

*“We switched from Oracle Database to IBM DB2 and cut our costs in half, while improving performance and reliability of business applications.”*

### JSC Rietumu Banka

- Moved from Oracle Database to IBM DB2
- Used “compatibility features”
- 3-30x faster query performance
- 200% improvement in data availability

9.7.1	SUB STRB,UDF Parameters: INOUT,FORALL/BULK COLLECT, Boolean, Conditional compilation, OCI Support	Increase compatibility
9.7.2	UDF Parameters: DEFAULT	Increase compatibility
9.7.2	Obfuscation	Enhancement
9.7.2	NCHAR, NVARCHAR, NCLOB	Increase compatibility
9.7.3	NUMBER Performance	Performance
9.7.3	Runtime “purity level” Enforcement	Increase compatibility
9.7.3	RATIO_TO_REPORT Function	Increase compatibility
9.7.3	RAISE_APPLICATION_ERROR	Increase compatibility
9.7.3	Small LOB Compare	Increase compatibility
9.7.4	Multi-action Trigger & Update Before Trigger	Increase compatibility
9.7.4	Autonomous Tx Improvements	Increase compatibility
9.7.4	LIKE Improvements, LISTAGG	Increase compatibility
9.7.4	ROW & ARRAY of ROW JDBC Support	Increase compatibility
9.7.5	Pro*C Support	Increase compatibility
9.7.5	Nested Complex Objects	Increase compatibility
10	Local Procedure Definitions	Increase compatibility
10	Local Type Definitions	Increase compatibility
10	PL/SQL Performance	Performance
10.1.2	SUBTYPE, PIPElined functions, .....	Increase compatibility
10.5	Rows Beyond 32K	Increase compatibility
10.5	Index on Expression	Increase compatibility
10.5	Compatible unique index semantics	Increase compatibility
10.5	@dblink syntax when referencing remote tables	Increase compatibility
10.5	INSTRB, INSTR2, INSTR4, LENGTHB, LENGTH2, LENGTH4, and SUBSTR4 functions	Increase compatibility

# Oracle Compatibility: Internationalization

## ▪ Character based semantics

- Allows applications written against single byte code page databases to work in multi-byte Unicode environment
- A character in Unicode can take between 1-byte (ASCII) and 4-bytes (excluding combining characters)
- Number of characters a column can hold is unclear
  - `CREATE TABLE emp(name VARCHAR(10));`
  - Can hold between 2 and 10 characters!
- Regular string functions are byte based
  - `SUBSTR('Äpfel', 1, 1) → Empty string`
  - `SUBSTR('Äpfel', 1, 2) → 'Ä'`
- String functions now treat character based strings correctly
 

```
SELECT english FROM monuments
WHERE SUBSTR(native, 1, 2) = '万里'
```

## ▪ Accommodate larger strings

- Allow tables with up to 1MB wide rows
 

```
CREATE TABLE emp(name VARCHAR(4000),
                  address VARCHAR(4000),
                  cv VARCHAR(32000))
```
- Allow large row GROUP BY and ORDER BY as long as key can sort

Native	English
Кремль	Kremlin
万里长城	Great Wall of China
ताज महल	Taj Mahal
Hoover Dam	Hoover Dam

# Oracle Compatibility: Additional Indexing

## ▪ Function-based indexes

- Searching for computed values in a table instead of using Generated Columns
- E.g. “Find employees without worrying about the case of their names”

- ```
CREATE INDEX emp_name ON emp(UPPER(name));
SELECT salary
  FROM emp
 WHERE UPPER(name) = 'MCKNIGHT';
```

| Name     | Salary | Manages |
|----------|--------|---------|
| McKnight | 50000  | Sales   |
| Miller   | 25000  | -       |
| Van Gogh | 45000  | Finance |
| Chan     | 37000  | -       |

## ▪ Indexes excluding NULL keys

- Enforce uniqueness only for non-NULL keys and exclude all NULL keys from Index
- Compress index for all-NULL keys
- Helps facilitate Oracle application migrations

- ```
CREATE UNIQUE INDEX emp_manages
  ON emp(manages) EXCLUDE NULL KEYS
```

## ▪ Random key indexes

- Avoid hot index page for incrementally issued keys

- ```
CREATE UNIQUE INDEX order_id ON order(id RANDOM);
```

---

# Oracle PL/SQL Compatibility

- **Create distinct type with weak type rules**
    - Removes limitation of existing distinct types not having weak typing
    - Optional check constraint
    - Optional NOT NULL constraint
    - Constraints enforced on assignment
  
  - **Pipelined table function**
    - Introduce a new PIPE statement which returns a row to caller, but continues at next statement if caller wants another row
    - Incrementally produce a result set for consumption on demand
  
  - **Ad-hoc federated table access**
    - Support ad-hoc reference to remote table using server in the identifier
      - Reach out to a table in a remote database
  
  - **Function library extensions**
    - Updates to various built-in functions for improved compatibility support
-

---

# Tools to Help You Move Easily from Oracle

- **Database Conversion Workbench (DCW)**

- Analyzes PL/SQL and database objects to quickly identify areas that may require manual intervention during conversion

<https://www.ibm.com/developerworks/community/groups/service/html/communityview?communityUuid=05901c97-75b2-47a1-9c32-25f748855913>

- Migrates Oracle data and schema objects to DB2
- Plug-in for Data Studio

- **DB2 Conversion Assessment Questionnaire**

- Used to gather information for accurately sizing the migration effort. Based on our experience performing hundreds of application migrations

- **IBM DB2 Workshop for Oracle Professionals**

- Designed especially for Oracle practitioners, provides hands-on proof of how easy it is to learn and apply their existing skills to DB2

- **Flexible Migration Services to fit your needs**

- We offer a variety of Service options, from helping your existing staff migrate applications on their own to turnkey migrated solutions

- **And much more!**





# IBM Mobile Database



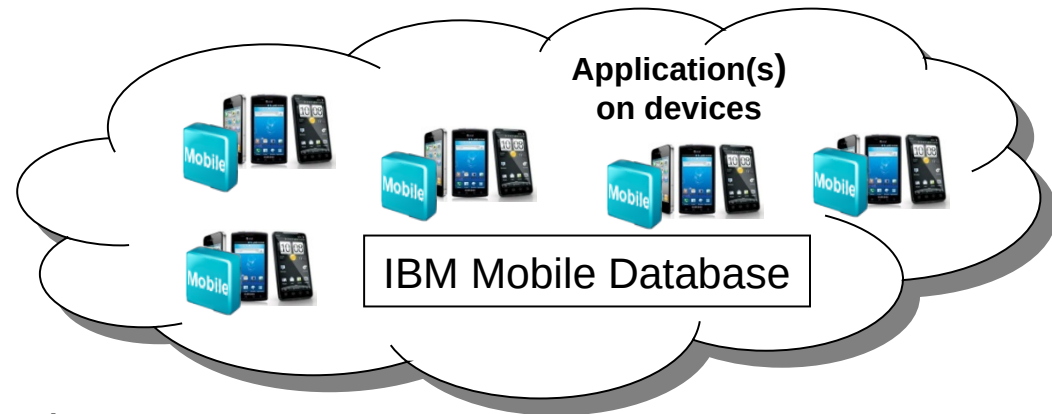
- **Full-featured, small footprint mobile data management solution**

- Persistent data
- Secure storage
- Synchronization with back-end databases

- **Available on Android**

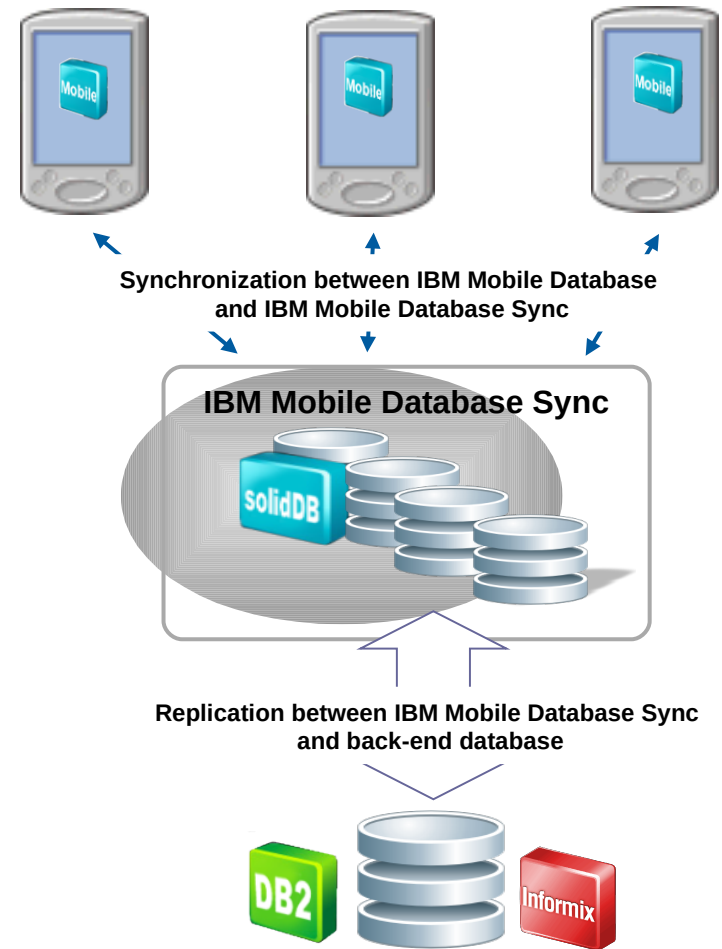
- IOS and Windows Mobile in development

- **Free to download from the web**



## Connectivity with Back-End Databases

- **IBM Mobile Database Sync gives rich synchronization capabilities for bidirectional communication between the IBM Mobile Database and enterprise databases**
  - IBM Mobile Database replicates data with IBM Mobile Database Sync
  - IBM Mobile Database Sync replicates data with back-end data servers
- **Back-end database can be IBM DB2 or IBM Informix**
- **Multiple solidDB systems can be used to scale the system for large number of devices**



# Background - What is JSON?

- **JavaScript Object Notation**
  - Serialized form of JavaScript Objects
    - Lightweight data interchange format
    - Specified in IETF RFC 4627
    - <http://www.JSON.org>
- **Lightweight text interchange**
  - Designed to be minimal, portable, textual, and subset of JavaScript
    - Only 6 kinds of values!
    - Easy to implement and easy to use
- **Replacing XML as the de facto data interchange format on the web**
  - Used to exchange data between programs written in all modern programming languages
- **Self-describing, easy to understand**
  - Text format, so readable by humans and machines
  - Language independent, most languages have features that map easily to JSON

```
{
  "firstName": "John",
  "lastName" : "Smith",
  "age"      : 25,
  "address"  :
  {
    "streetAddress": "21 2nd Street",
    "city"         : "New York",
    "state"        : "NY",
    "postalCode"  : "10021"
  },
  "phoneNumber":
  [
    {
      "type" : "home",
      "number": "212 555-1234"
    },
    {
      "type" : "fax",
      "number": "646 555-4567"
    }
  ]
}
```

*“Less is better: less we need to agree upon to interoperate, the more easily we interoperate”*

JavaScript: The Good Parts, O'Reilly

# DB2 Technology Preview - JSON

- **Combine data from systems of engagement with traditional data in same DB2 database**
  - Best of both worlds
  - Simplicity and agility of JSON + enterprise strengths of DB2
  
- **Store data from web/mobile apps in it's native form**
  - New web applications use JSON for storing and exchanging information
  - It is also the preferred data format for mobile application backends
  
- **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



# DB2 10.5 for zLinux – What's changed ?

- **Administration**
  - HADR, database manager configuration, registry variables, ...
- **Database setup**
  - Database configuration,
- **Application development**
  - System catalog views, build-in functions, global variables, ...
- **DB2 command and SQL statements**
- **Deprecated and discontinued functionality**

Detail information in DB2 Information Center:

<http://publib.boulder.ibm.com/infocenter/db2luw/v10r5/index.jsp?topic=%2Fcom.ibm.db2.luw.wn.doc%2Fdoc%2Fc0060311.1>

# So What Is DB2 for zLinux? (and that is it not)

- **DB2 LUW for zLinux is**
  - “JUST” DB2 LUW
  - 99.9%+ the same code
  - 99.9%+ the same look/feel
  - 99.9%+ the same skills
  
- **DB2 LUW for zLinux has 98%+ Oracle Compatibility**
  - PL/SQL
  - Datatypes
  
- **DB2 LUW is not DB2 for zOS, but**
  - Common client infrastructure
    - JCC/ODBC/CLI
    - Perl/php/python/Ruby/...
  - Optim Data Studio tools (Data Studio, OPM, OPWT)
  - Exceptionally similar DDL/DML/DCL

# DB2 for z/OS vs DB2 for LUW

(not exhaustive, features may be missing)



DB2 10 for z/OS and DB2 10.1 Linux, Unix & Windows

z/OS

{ Multi-row INSERT, FETCH & multi-row cursor UPDATE, Dynamic Scrollable Cursors, GET DIAGNOSTICS, Enhanced UNICODE SQL, join across encoding schemes, IS NOT DISTINCT FROM, VARBINARY, FETCH CONTINUE, MERGE, SELECT from MERGE, routine versioning, timestamps w/timezone

common

{ Inner and Outer Joins, Table Expressions, Subqueries, GROUP BY, Complex Correlation, Global Temporary Tables, CASE, 100+ Built-in Functions including SQL/XML, Limited Fetch, Insensitive Scroll Cursors, UNION Everywhere, MIN/MAX Single Index, Self Referencing Updates with Subqueries, Sort Avoidance for ORDER BY, and Row Expressions, 2M Statement Length, GROUP BY Expression, Sequences, Scalar Fullselect, Materialized Query Tables, Common Table Expressions, Recursive SQL, CURRENT PACKAGE PATH, VOLATILE Tables, Star Join Sparse Index, Qualified Column names, Multiple DISTINCT clauses, ON COMMIT DROP, Transparent ROWID Column, Call from trigger, statement isolation, FOR READ ONLY KEEP UPDATE LOCKS, SET CURRENT SCHEMA, Client special registers, long SQL object names, SELECT from INSERT, UPDATE or DELETE, INSTEAD OF TRIGGER, Native SQL Procedure Language, BIGINT, file reference variables, XML, FETCH FIRST & ORDER BY in subselect & fullselect, caseless comparisons, INTERSECT, EXCEPT, not logged tables, OmniFind, spatial, range partitions, data compression, session variables, DECIMAL FLOAT, optimistic locking, ROLE, TRUNCATE, index & XML compression, created temps, inline LOB, administrative privileges, implicit cast, date/time changes, currently committed, moving sum & average, index include columns, row and column access control, time travel query

LUW

{ Updateable UNION in Views, GROUPING SETS, ROLLUP, CUBE, more Built-in Functions, SET CURRENT ISOLATION, multi-site join, MERGE, MDC, XQuery, XML enhancements, additional data type (array, row, cursor), global variables, even more vendor syntax, temp table compression, MODULES

# Strong Value Scenarios

## ▪ Consolidation

- Bring many small, typically lightly loaded databases together into a highly consolidated and virtualized environment under z/VM

## ▪ Migrate unhappy Oracle on zLinux customers

- Oracle may be pushing them to Exadata
- Oracle may be charging outrageous prices
- Migration is typically relatively easy/trivial
  - Especially if using a packaged app that already supports DB2

## ▪ Use DB2 as the persistent data store inside another SWG solution

- Websphere Portal Server
- Cognos 10 content server
- ... many more



## Best Practices (System config)

- **Use a modern System z machine**
  - The System z processor has caught up a lot on a per core (IFL) basis
  
- **If you want to use z/VM (most do) remember**
  - z/VM is not “free” (resource-wise) – especially wrt. memory
    - The more IFLs the more memory you need to reserve for z/VM
  - Ensure you have current maintenance for z/VM
    - Every test we do “finds” interesting things.
  
- **If your database does disk I/O (most do) remember**
  - These days you care less about storage space
    - You care about storage performance (random IOPS, sequential MB/sec)
  - Avoid ECKD disk (it is not a good fit for DB2), use SCSI disks
  
- **Use a current Linux version**
  - SLES 11 or RHEL6
  - Consult with the zLinux performance team on specifics

## Best Practices (DB2 software/sizing)

- **Use current DB2 LUW software**
  - I know the tendency on zOS is to use n-1 VERSIONS
  - Linux is different
  
- **DB2 10.1 has been well-received (across platforms)**
  - Many customers have tested with DB2 10.1 FP1 in 2012
  - Many production deployments of DB2 10 in 1H2013
  - DB2 10.1 FP2 is already GA
  
- **Don't overthink the sizing. First order approximations**
  - DB2 for zOS = DB2 LUW
  - DB2 LUW = Oracle
  - DB2 LUW on Power7 = DB2 LUW on zLinux on a z196
  
- **KISS principle**
  - Resist the urge to over-specify all possible options
    - Just because they exist doesn't mean you have to use them

## Best Practices (DB2 configuration)

- **Use DB2 automatic storage**
  - A modest number of storage paths that map sensibly to the SAN storage
- **Minimize the use of distinct page sizes**
  - 8K is a good default (no more than 2)
  - Specify the smaller one at create database time for the catalog
- **Minimize the number of bufferpools**
  - 1 for data/index, 1 for temp is usually sufficient
- **Go ahead and use plenty of tablespaces**
  - Often helpful to split data/index/LOB (esp LOB) into separate tablespaces
  - No need to put EACH table in its own tablespace
- **Use the Autonomics**
  - Autoconfigure & STMM
    - But give DB2 a budget via instance\_memory or database\_memory
- **Compression**
  - Turn on compression at create table time for top ~50 large tables

## Best Practices (Tools)

- **There is now a strong tools portfolio**
  - Many are free
  - Even more are included in Advanced Enterprise Server Edition (AESE)
  
- **Optim Data Studio**
  - Database Administration
  - Developer
  
- **Optim Performance Manager (OPM)**
  - With Extended Insight for application server environments (e.g. WAS)
  
- **Optim Query Workload Tuner (OQWT)**
  - Integration with OPM
  - Next generation of “DB2 design advisor”
  
- **Optim Capture/Replay**
  - Realistic Pre-production

---

## Summary

- **DB2 LUW is a mature product on zLinux**
  - Many years of development/testing/deployment
- **DB2 LUW on zLinux is full function**
  - Nearly every capability is available on zLinux
- **DB2 LUW on zLinux is high performance**
  - Regular testing on latest System z hardware
- **DB2 10.5 is available on zLinux**
- **100's of satisfied customers**

THANK

YOU