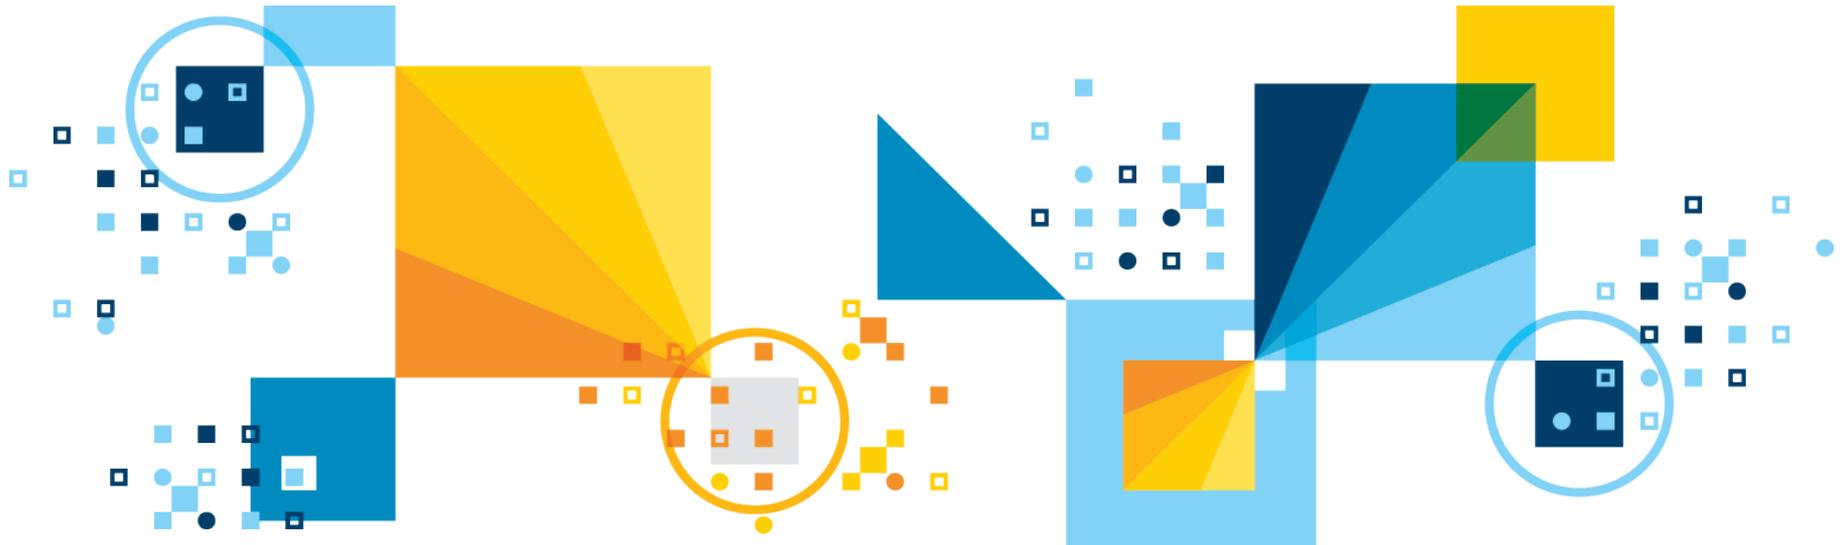


Jessica Rockwood & Roger E. Sanders

DB2Night Show #195

June 22, 2017

# Db2 11.1 Update Packaging, Functionality, Roadmap



# Please note

- © IBM Corporation 2017. All rights reserved. U.S. Government Users Restricted Rights - use, duplication, or disclosure restricted by GSA ADP Schedule Contract with IBM Corporation.
- IBM, the IBM logo, ibm.com and DB2 are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or TM), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at “Copyright and trademark information at : [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).
- The information contained in this presentation is provided for informational purpose only. While efforts were made to verify the completeness and accuracy of the information contained in this presentation, it is provided “as is” without warranty of any kind, expressed or implied. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this presentation or any other documentation.
- 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. Nothing contained in this presentation is intended to, nor shall have the effect of, creating any warranties or representations from IBM (or its suppliers or licensors), or altering the terms and conditions of any agreement or license governing the use of IBM products and/or software.
- Any statements of performance are 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 multi-programming 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.
- IBM’s statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM’s sole discretion. The development, release, and timing of any future features or functionality described for our products remains at our 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.

# Agenda

- **Updates to the Db2 Brand**
- **11.1.2.2 Licensing and Packaging**
- **Download and Go**
- **11.1.2.2 Content Details**
- **Roadmap Overview and Additional Resources**

“Data is the world’s next natural resource. The value of the data and extracting the information and insights will change how we make every decision.”

Ginni Rometty



# Agenda

- **Updates to the Db2 Brand**
- 11.1.2.2 Licensing and Packaging
- Download and Go
- 11.1.2.2 Content Details
- Roadmap Overview and Additional Resources

# Updates to the Db2 Brand

## 1. Db2 the brand

- IBM is *renaming DB2* ... modifying the casing to result in **Db2**
- The new name will be rolled out on webpages, Knowledge Centers, product names, etc. over the coming weeks and months
- Internal codes will remain DB2 to prevent an impact on applications that use this information

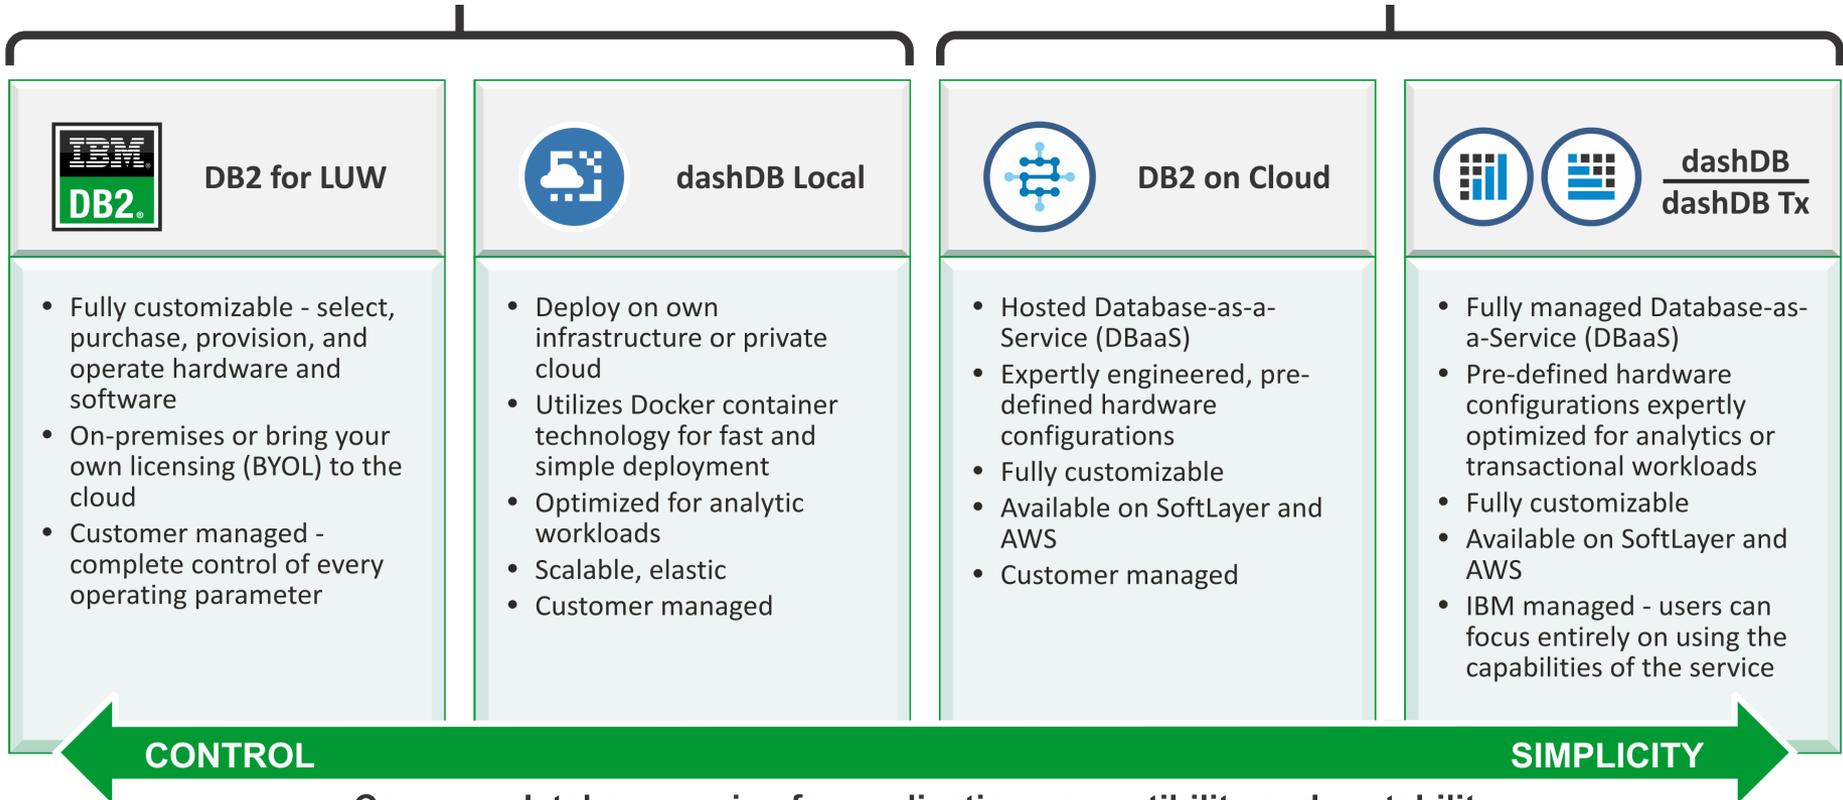
## 2. Db2 the product

- Previously known as *DB2 for Linux, Unix and Windows (LUW)*, the Db2 which runs on distributed platforms is now *simply Db2*
- The identification of platform remains for Db2 for z/OS and Db2 for i

# Product names prior to Db2 11.1.2.2

## On-premises or Private Cloud

## Public Cloud



**CONTROL**

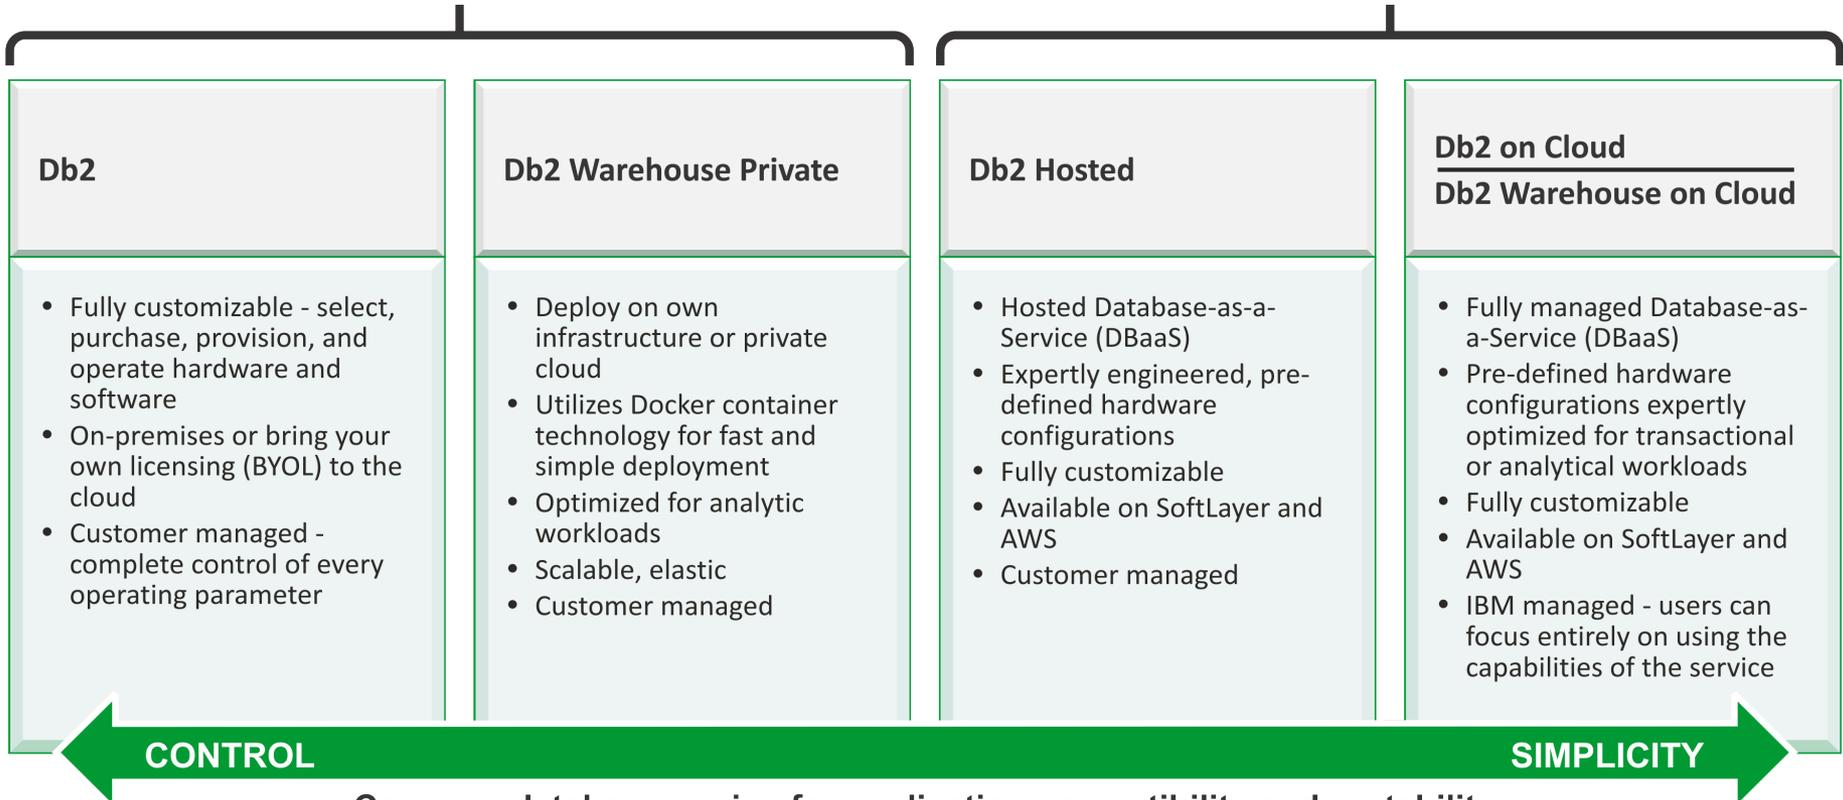
**SIMPLICITY**

Common database engine for application compatibility and portability

# Product names starting with Db2 11.1.2.2

## On-premises or Private Cloud

## Public Cloud



Common database engine for application compatibility and portability

# Agenda

- Updates to the Db2 Brand
- **11.1.2.2 Licensing and Packaging**
- 11.1.2.2 Content Details
- Download and Go
- Roadmap Overview and Additional Resources

# Current Development Ecosystem

## ▪ Current: Db2 Express-C:

- Perfect for developers and ISVs as it allows them to develop and deploy applications at no charge, with some environment and functional limitations
  - 2 cores, 16GB of memory, 15TB of storage
  - No pureScale, no DPF, no BLU Acceleration, no compression, etc.
- Great for academic communities, which in turn helps make more Db2 skills available in the marketplace (with those skills limited based on Db2 capabilities of Express-C)
- Available on Linux and Windows x64 operating environments
- Not the common package image as ALL other Db2 editions
  - can not be converted to another edition (via db2licm)
- Download Analysis: ~70,000 downloads / year



## ▪ Db2 Developer Edition

- For application developers to design, build, and prototype applications for deployment
- Includes all the Db2 server editions allowing for solutions to be built with the latest data server technologies
  - Cannot be used for production systems
- An Authorized User license is required for each person with access
  - Only 1 license required per person – access to any and every Db2 Developer

# Expand Db2 Ecosystem/Skills

## ▪ New Db2 Developer-C Edition

- Freely downloadable just like Express-C
- Fully functional Db2 database server
  - pureScale and DPF deployments
  - Compression and BLU Acceleration

## ▪ Use Limitations

- For development and non-production only
- **Unsupported edition (non-warranted)**
- Environment limited usage:
  - 4 cores, 16GB of memory
  - 100GB of data in user tablespaces

## ▪ Functional Differences from Db2 Developer Edition

- Excludes some “Supporting Programs”
  - Cognos, IDA, WebSphere AS / MQ, DSM Enterprise



# Db2 Developer-C Supporting Programs

## ▪ Supporting Programs included with Db2 Developer-C

- IBM Database Add-ins for Visual Studio
- IBM Data Server Drivers (all)
- IBM Tivoli System Automation for Multiplatforms
- IBM Enterprise Content Management Text Search
- IBM Security Directory Server
- IBM Global Security Toolkit
- IBM Spectrum Scale
- IBM DB2 Connect Enterprise Edition
- IBM InfoSphere Data Replication (CDC component)
- IBM Data Studio



## ▪ Db2 Advanced/Developer supporting programs NOT included

- IBM Data Server Manager Enterprise Edition
- IBM InfoSphere Data Architect
- IBM Cognos Analytics
- IBM WebSphere Application Server
- IBM WebSphere MQ

# *Non-Production* License Metric for Db2 Direct Editions

- **Db2 Non-Production parts** for all Db2 Direct Standard/Advanced
  - Non-Production use only
  - Example uses include dev, test, qa, benchmarking, etc
  - New *license metric*
- **Pricing**
  - **15% discount** on Direct Standard / Advanced
- **Fully Supported**
  - Not a separate edition, simply discounted price metrics on existing editions
- **Proposed Change to License Agreements**
  - Non-Production Limitation

If the Program is designated as "Non-Production", the Program can only be deployed as part of the Licensee's internal development and test environment for internal non-production activities, including but not limited to testing, performance tuning, fault diagnosis, internal benchmarking, staging, quality assurance activity and/or developing internally used additions or extensions to the Program using published application programming interfaces. Licensee is not authorized to use any part of the Program for any other purposes without acquiring the appropriate production entitlements.



## End of Service for Db2 Versions 9.7 and 10.1

- **End of Service for both Db2 Version 9.7 and 10.1** was announced in 2016 in conjunction with the announcement of Db2 Version 11.1
  - Effective End of Service date of **September 30<sup>th</sup>, 2017**
  - Db2 11.1 supports **direct upgrade from Db2 Versions 9.7, 10.1, and 10.5**
  
- **NEW: Continuing Support**
  - <https://www-01.ibm.com/support/docview.wss?uid=swg22002945>
  - Maintaining S&S entitles customers to continue to open PMRs
    - Investigate issue, identify workaround, etc.
  - If solution *already exists (built on customers platform)*, fix can be provided
    - **No** new fixes or new builds will be created and provided
  
- **What if a customer cannot move off 9.7 or 10.1 before EOS date?**
  - **Extended support contracts** can be purchased for those customers who require a longer time to upgrade (preferably to 11.1)
  - IBM Db2 End of Service dates are **not applicable to SAP/Db2 customers**

# Summary

- **Express-C and Developer Editions**
  - No changes or impacts to clients using these editions
- **Developer-C Edition**
  - Fully functional Db2 for non-production only
  - No support
  - Builds Db2 skills on Advanced technologies
  - Limits on cores, memory, and data storage
  - Easy to convert into a fully licensed installation
- **Non-production licensing options**
  - Discounted licensing for all metrics on every Db2 edition
  - Fully supported
- **Continuing support**
  - PMRs can continue to be opened
  - No new fixes or special builds

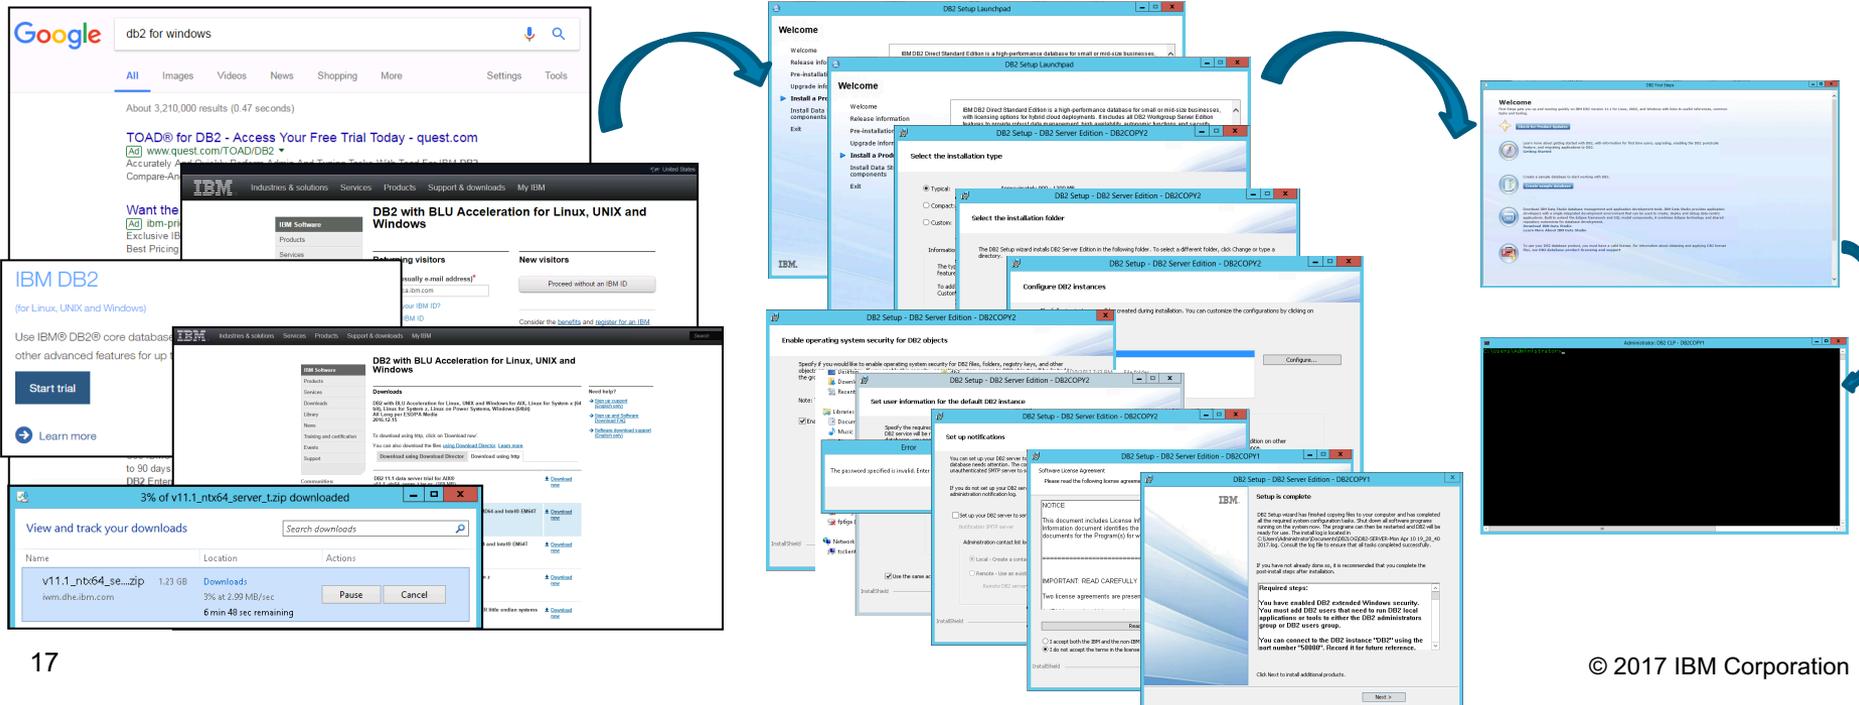


# Agenda

- Updates to the Db2 Brand
- 11.1.2.2 Licensing and Packaging
- **Download and Go**
- 11.1.2.2 Content Details
- Roadmap Overview and Additional Resources

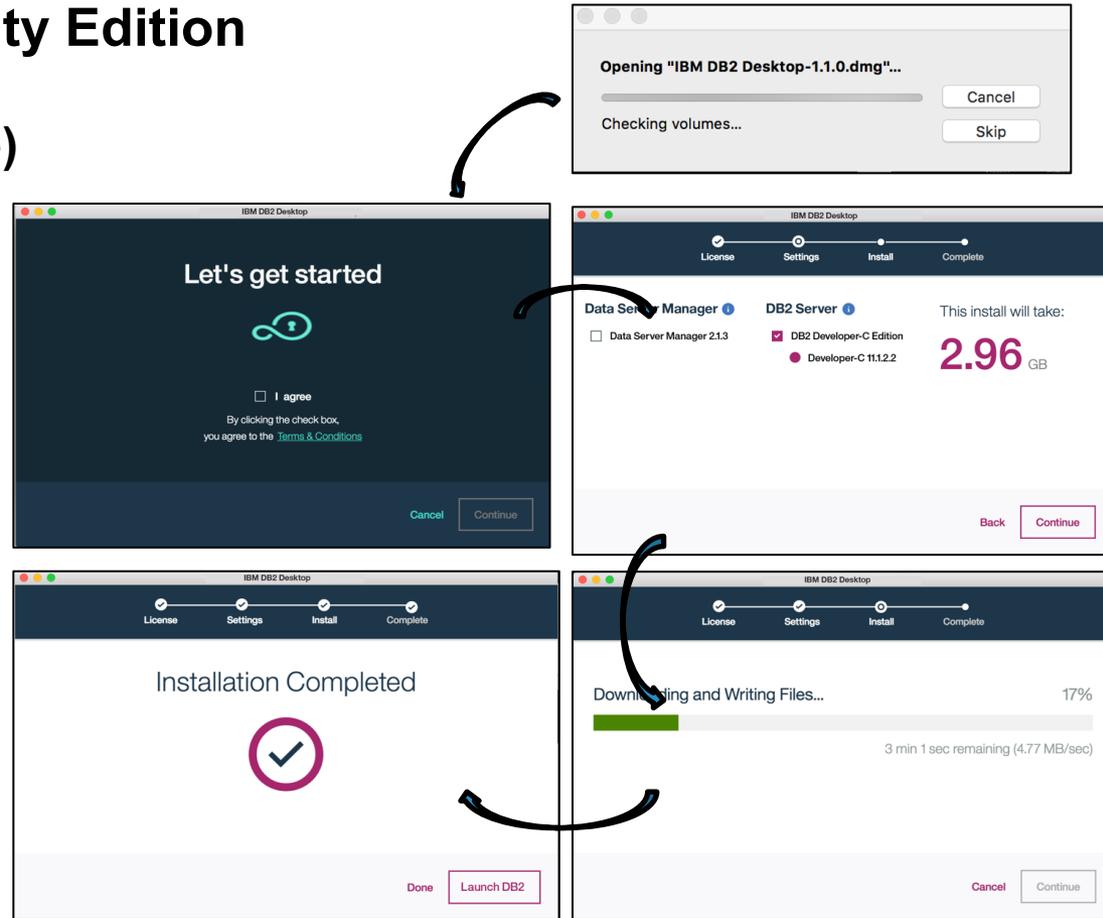
# Getting Db2 software

1. Time-consuming google search, ibm.com navigation, IBM id + password, license acceptance – 10+ minutes to download
2. Untar archive, find executable and install with outdated interface, 10+ screens with significant input required – 10+ minutes to install
3. No default graphical interface included with deployment, post-install limited “Db2 First Steps” and command line, optional DSM download



# Announcing the new “Download and Go” experience! IBM Db2 Developer Community Edition

- **Small download (Docker image)**  
**less than 15 minutes** from  
**“I want” to “I got”**
- **Platform specific executable**
  - Mac, Windows, Linux
- **From 10+ screens and 10+ mandatory input fields to 3 clicks and userID/pwd input only**
- **Download package includes**
  - Db2 Developer-C 11.1.2.2
  - Data Server Manager 2.1.4
  - Data Studio (optional)
- **Same limitations as Developer-C**
  - Development and **non-production only**
  - Core, memory, and storage **limitations**
  - **Unsupported** edition (non-warranted)



## Download and Go website

**<https://www.ibm.com/us-en/marketplace/ibm-db2-direct-and-developer-editions>**

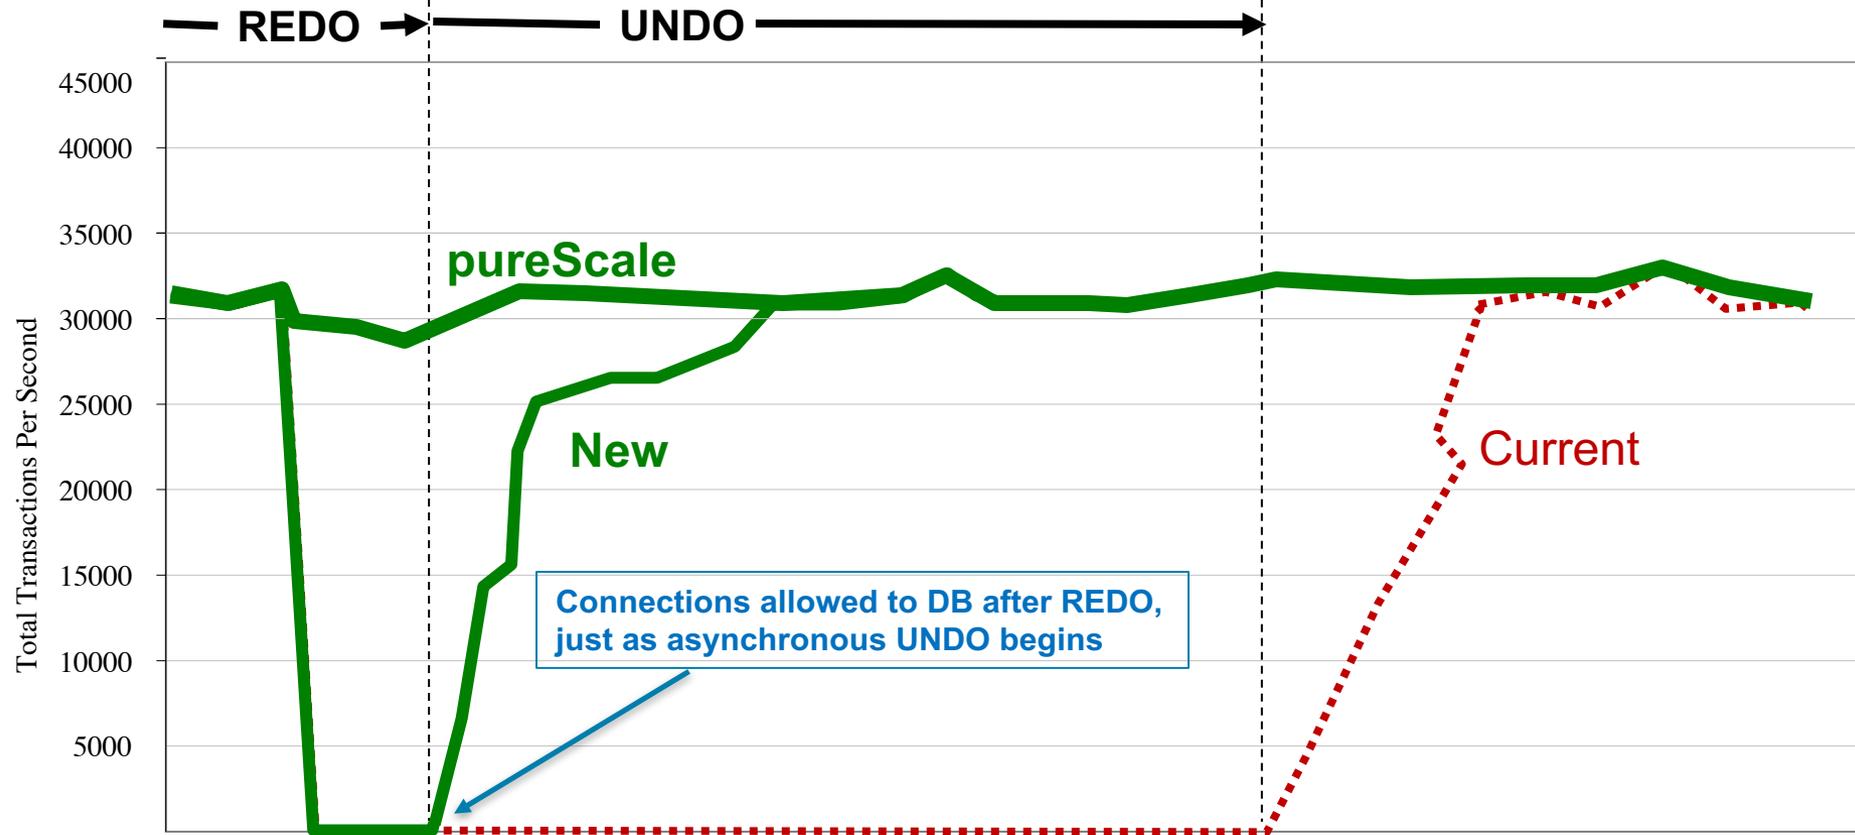
# Agenda

- Updates to the Db2 Brand
- 11.1.2.2 Licensing and Packaging
- Download and Go
- **11.1.2.2 Content Details**
- Roadmap Overview and Additional Resources

## Online Crash Recovery (Async UNDO)

- **Prior to 11.1.2.2, during Crash Recovery (or HADR Takeover by Force), connections are not allowed to the database until all recovery is complete**
  - A REDO phase to redo all transaction activity to the end of log
  - An UNDO phase to rollback transactions that had not committed
  
- **In Db2 11.1.2.2, we **allow connections and activity into the database after the REDO phase and while the UNDO phase is executing****
  - This was a technical preview in Db2 11.1.1.1

# Online Crash Recovery (Async UNDO)



(\*) Some data may be locked while UNDO proceeds concurrently

# Recoverability Enhancements

## ▪ Improved Transaction Rollback Performance

- Additional Db2 process now dedicated to maintaining recovery starting point
- Internal test scenario of 4 concurrent applications inserting & rolling back showed **5-6X faster rollback performance**

## ▪ Faster pureScale Member Crash Recovery

- Improved hashing algorithm and increased parallelism in global lock manager recovery design
  - Reduces the recovery window
- Benefits all pureScale deployments
  - Particularly beneficial for large clusters configured with huge lock list, global lock manager, and global buffer pool

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.

# Availability enhancements

## ▪ HADR Tablespace Recovery

- If the standby database has a tablespace in an invalid or error state, a TAKEOVER by the Standby will not be fully available to applications
- Prior to Db2 11.1.2.2, the entire Standby database needed to be recovered
- In 11.1.2.2, the tablespace can be recovered online without a full database refresh
- Tablespace recovery been retrofitted into Db2 10.5fp9+
- Details: <http://www-01.ibm.com/support/docview.wss?uid=swg21993389>

## ▪ Restore REBUILD in pureScale

- Prior to 11.1.2.2, pureScale required a database backup in order to perform a database recovery
- Functionality available today in non-pureScale deployments
- In 11.1.2.2, pureScale can now recover a database from a set of tablespace backups

# Monitoring HADR Tablespace Status

## ▪ Tablespace Error State

- When a tablespace is in an invalid or error state on the Standby database, the **HADR\_FLAGS** field will display the value **STANDBY\_TABLESPACE\_ERROR**

## ▪ Monitoring with db2pd

- The **HADR\_FLAGS** field can be monitored by using the **db2pd -hadr** command on the Primary or Standby database

```
db2pd -db HADRDB1 -hadr
```

```
... HADR_FLAGS = STANDBY_TABLESPACE_ERROR TCP_PROTOCOL ...
```

## ▪ Monitoring with Table Function

- The **MON\_GET\_HADR()** table function will display the current status on either the Primary database or Standby database with Reads on Standby enabled

```
SELECT STANDBY_ID, HADR_FLAGS FROM
       TABLE (MON_GET_HADR (NULL) )
```

```
STANDBY_ID
```

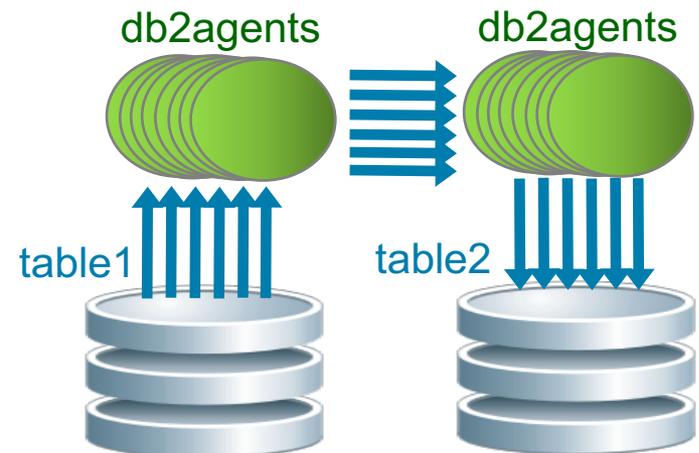
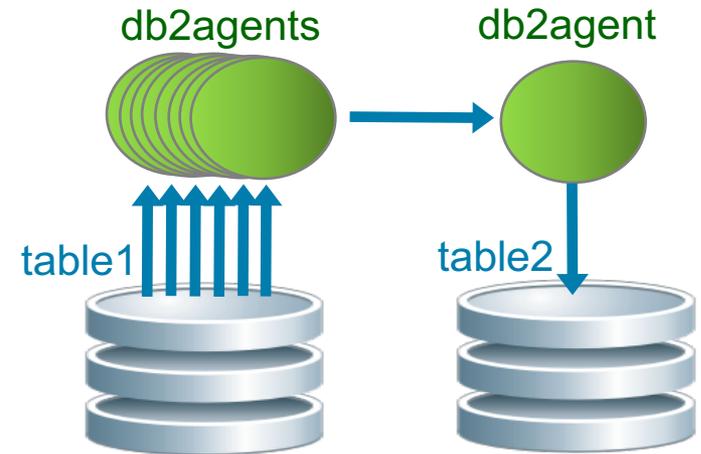
```
-----
```

1	STANDBY_TABLESPACE_ERROR TCP_PROTOCOL
---	---------------------------------------

# Multi-Core Parallelism – BLU INSERT

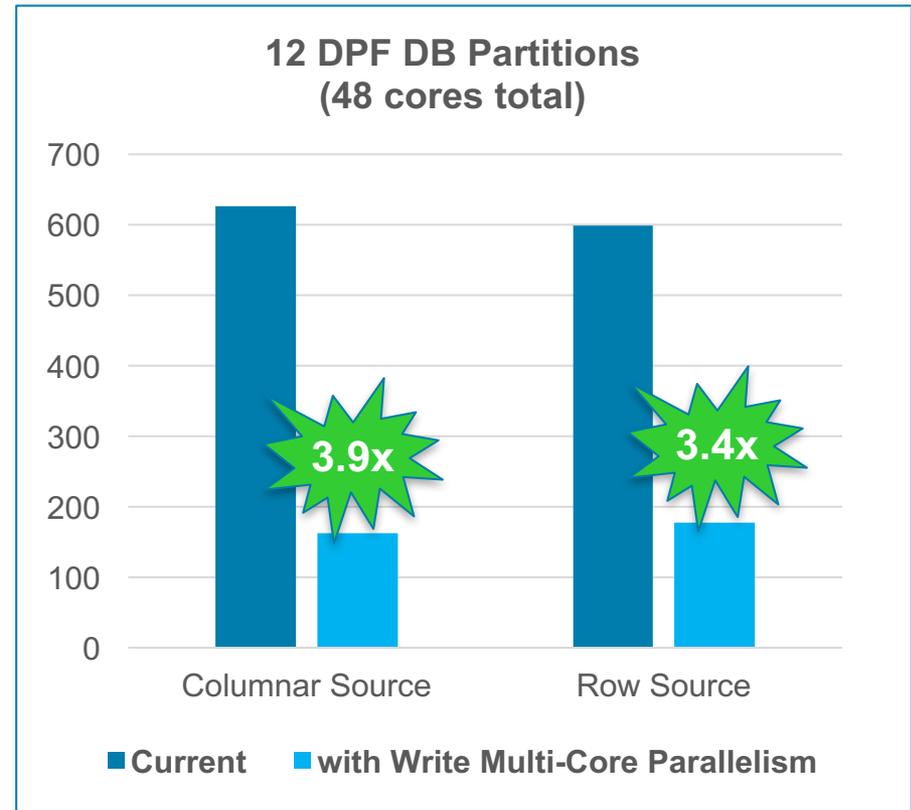
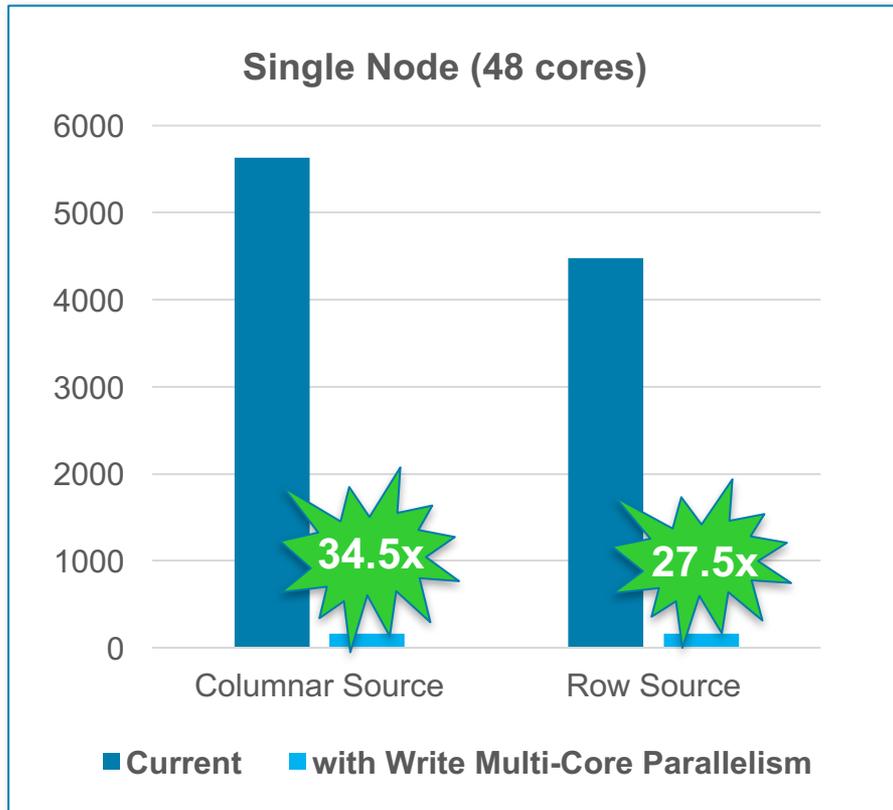
- **Db2 has exceptional multi-core scalability for queries:**
  - Combined MPP and SMP parallelism
  - Includes columnar and row-based tables
  - Applies to sub-select of INSERT
  
- **New Multi-Core Parallelism for INSERT**
  - Includes a variety of data population methods (i.e., INSERT, INGEST, etc.)
  - Extends to index maintenance
  - Also applicable to MPP deployments
  
- **Very significant reduction in time for:**
  - ETL/ELT batch jobs
  - Data population/ingest jobs
  - Move/copy data from one table into another

```
INSERT INTO table2
SELECT * FROM table1
```



# Multi-Core INSERT Parallelism – Example Benefits

▪ Elapsed time (seconds) for: `INSERT INTO table2 SELECT * FROM table1`



- 100 Million Rows
- 47GB
- DPF tests used 12 logical DB partitions, & co-located tables
- Same 48 core server used for single-node and DPF tests

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 27 configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

# Performance Enhancements (1 of 2)

## ▪ Adaptive Compression

- Decompression efficiency improved with lab tests showing ~30% cpu savings

## ▪ DECFLOAT

- Hardware acceleration of DECFLOAT operators on supported platforms

## ▪ COUNT/BIG\_COUNT/SUM

- Optimized operators in row and column based systems

## ▪ TRUNC, MOD, ROUND

- Constants in an SQL statement are evaluated at compile time rather than executed continuously at run time

```
SELECT * FROM EMPLOYEE WHERE SALARY > TRUNC(50000.8) ...
```

- The `TRUNC(50000.8)` would be evaluated before running the query

## ▪ RENAME TABLESPACE

- Catalog cache maintenance takes time and impacts concurrency
- Changes made internally to optimize memory usage in catalog cache
- Command runs much more efficiently

## Performance Enhancements (2 of 2)

- **Query performance: BLU synopsis exploitation for LIKE predicates**

- Data skipping capabilities leveraged by queries using LIKE predicates

- **Improved text search performance for queries containing OR predicates used with the text search functions**

- Queries are converted (query rewrite) to use UNION DISTINCT
- Example: books table contains 1000 rows

```
SELECT title FROM books WHERE CONTAINS(title, 'abc')=1
                                OR CONTAINS(title, 'def')=1
                                OR CONTAINS(title, 'ghi')=1
```

- Previous design calls TS server *for each row for each value* – 3000 times
- New design calls TS server *once for each value*– 3 times

# SAP Performance Enhancements

## ▪ Redundant Function Elimination

- Eliminate functions in SQL statements that do not effect the outcome

```
SELECT * FROM EMPLOYEE WHERE RTRIM(RTRIM(NAME)) = 'HAAS'
```

## ▪ Reduced Overhead in Unfenced C-UDFs Calls

## ▪ Complex CASE Statement Support (Increased Size)

## ▪ SAP BW Use of Parallel Insert (Up to 11x faster ETL processing)

## ▪ SAP DBA Cockpit exploitation of DB2\_GET\_INSTANCE\_INFO for Windows

## ▪ Extend CHANGE HISTORY event monitor to include the following:

- exclude / include logs, number of sessions, parallelism, size of the buffers
- number of buffers, UTIL\_IMPACT\_PRIORITY, usage of deduplication
- usage of compression with which compression library and which options
- usage of encryption with which encryption library and which options
- snapshot backup with which vendor library and which options
- scripted snapshot with which script

# Package Cache Efficiency

- **Current statement concentrator has two options:**
  - OFF: Statement Concentrator is off
  - LITERAL: Literals are removed from the SQL and compiled with parameter markers for more efficient package cache use
- **Using LITERAL does not work well when the SQL statement includes optimization information or comments**
  - SQL statements could be the same but the comments make them look different in the package cache
- **Two new options for the STMT\_CONC database configuration parameter**
  - COMMENTS: concentrate statements based on comments. All comments in the SQL statements would be removed.
  - COMM\_LIT: concentrate statements based on comments and literals in the SQL statement text.
- **Syntax:**  
`db2 update db cfg using STMT_CONC COMMENTS`

# Package Cache Efficiency

## ▪ Example 1:

```
SELECT EMPNO FROM EMPLOYEE --OPTLEVEL (5) --QUERY_DEGREE (1)
```

```
SELECT EMPNO FROM EMPLOYEE --OPTLEVEL (6) --QUERY_DEGREE (1)
```

- LITERAL: Each statement would be considered different in the package cache
- COMMENTS: Each statement is considered the same and we can reuse the package cache

## ▪ Example 2

```
SELECT EMPNO FROM EMPLOYEE WHERE EMPNO = 5 --OPTLEVEL (5)
```

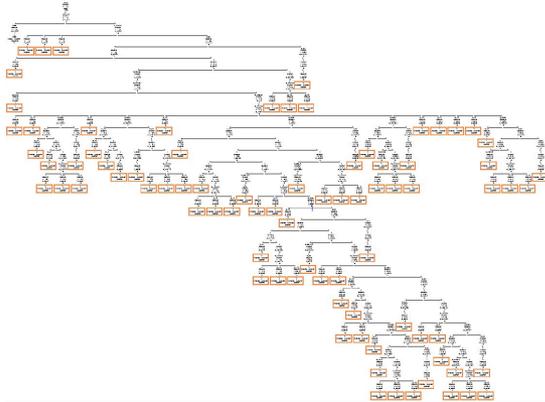
```
SELECT EMPNO FROM EMPLOYEE WHERE EMPNO = 6 --OPTLEVEL (6)
```

- COMMENTS: Each statement would be considered different in the package cache
- COMM\_LIT: Each statement is considered the same because the literals are replaced with parameter markers and we can reuse the package cache

## ▪ **MON\_GET\_PKG\_CACHE\_STMT** now includes **STMT\_COMMENTS** to view comment information in original SQL statement

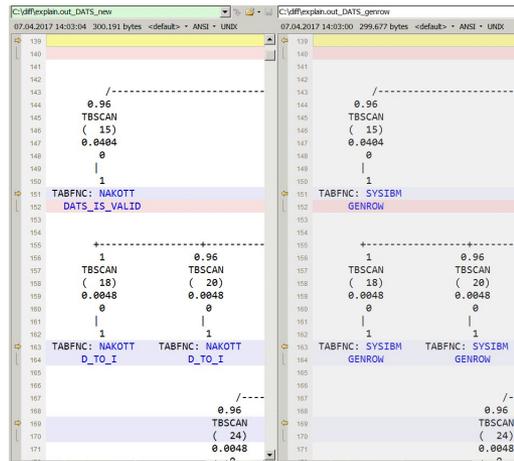
# Updated Explain

- User-defined Table Functions in Explain Plans are listed as GENROW



Every Yellow box represents a table function – But which one?

- Explain Information Updated to include Function Name and Fragment Details
  - Easier to determine which function was involved in a branch of the explain plan





## Other Enhancements

### ▪ **ADMIN\_MOVE\_TABLE updates**

- Transfer original ROW CHANGE TIMESTAMP content under ALLOW\_READ\_ACCESS when using COPY\_USE\_LOAD
- Optimize statistics during AMT

### ▪ **DB2\_AVOID\_LOCK\_ESCALATION registry variable**

- Provides ability to prevent lock escalation when either MAXLOCKS or LOCKLIST thresholds are hit
- Application will return an SQL0912N error

### ▪ **COL\_TABLE\_SERIALIZE lock details externalized**

- Internal lock for serializing IUD activity on column organized tables
- Now externalized in all lock monitoring (monitor, snapshot, db2pd, etc)

### ▪ **Improved logging for **db2chkupgrade** to **pinpoint invalid objects****

- Continuing to improve our messages for better self-help
- Identify names of the invalid object to avoid customers to search for them

### ▪ **New sub-options to db2trc for –cli : **-clitracelevel** and **-pattern****

- [https://www.ibm.com/support/knowledgecenter/en/SSEPGG\\_11.1.0/com.ibm.db2.luw.admin.cmd.doc/doc/r0002027.html](https://www.ibm.com/support/knowledgecenter/en/SSEPGG_11.1.0/com.ibm.db2.luw.admin.cmd.doc/doc/r0002027.html)

# XML Enhancements

## ▪ Implicit Names for XML Query

- Users can refer to a document without assigning a name. The XPATH expression won't need a document name at the beginning of the XPATH
- Compatible with syntax supported by Db2 for z/OS

### Required syntax prior to 11.1.2.2

1. SELECT XMLQUERY('\$i/root/customerinfo[@Cid=1003]' passing INFO **as "i"**) FROM customer;
2. SELECT T.\* from xs.ORDERS, XMLTABLE('\$i/order[orderedby=\$name]' passing xmlcol **as "i"**, 'John Doe' as "name" columns name VARCHAR(30) path 'orderedby', item INT path 'items/item/@id', quantity INT path 'items/item/quantity') AS T;

### New and simpler syntax supported

1. SELECT XMLQUERY('/root/customerinfo[@Cid=1003]' passing INFO) FROM customer
2. SELECT T.\* from xs.ORDERS, XMLTABLE('/order[orderedby=\$name]' passing xmlcol, 'John Doe' as "name" columns name VARCHAR(30) path 'orderedby', item INT path 'items/item/@id', quantity INT path 'items/item/quantity') AS T;

# Federation Enhancements

## ▪ Support for PostgreSQL and MySQL

- No longer need to use the generic ODBC data source and wrapper
- Direct knowledge of these newly supported data sources
- Two simple DDL statements to query open source databases

```
create server POST type postgresql version 6.0 options (host  
'9.112.250.82', dbname 'testdb', password 'Y', pushdown 'Y');
```

```
create user mapping for user server POST options(REMOTE_AUTHID  
'iidev6', REMOTE_PASSWORD 'password');
```

## ▪ Increased security with the ability to configure Secure Socket Layer (SSL) to connect to Db2 family data sources

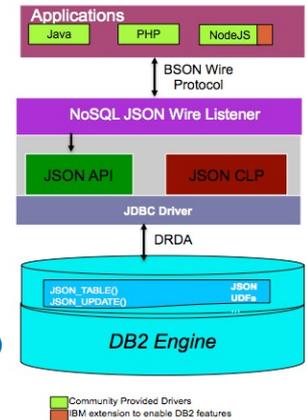
- [https://www.ibm.com/support/knowledgecenter/en/SSEPGG\\_11.1.0/com.ibm.db2.luw.admin.sec.doc/doc/t0070301.html](https://www.ibm.com/support/knowledgecenter/en/SSEPGG_11.1.0/com.ibm.db2.luw.admin.sec.doc/doc/t0070301.html)

## ▪ Integrated setup/configuration for ODBC drivers on PPC64LE

- Previously available on all platforms except PPC64LE

# JSON support

- **Initial JSON support in Db2 10.5 FP1 consisted of:**
  - IBM NoSQL Wire Listener, JSON Java APIs, JSON command line
- **Db2 11.1.2.2 will add full support for direct SQL access to underlying JSON interfaces used by externals above**
  - Integrated installation/configuration during database creation or migration update
    - No need to catalog JSON functions manually
  - High-level documentation on their existence similar to what Db2 for z/OS already provides
    - JSON2BSON, BSON2JSON are catalogued as part of SYSIBM schema
    - JSON\_VAL is function built into the SQL syntax
    - All other JSON functions are part of the SYSTOOLS schema
  - Overview of all functions and how to use them can be found in Db2 11 ebook at: <http://ibm.box.com/v/DB2v11eBook>



## 11.1.2.2 SQL functions

Schema	Name	Comments
SYSTOOL	BSON2JSON	Convert BSON formatted document into JSON strings
SYSTOOL	BSON_VALIDATE	Checks to make sure that a BSON field in a BLOB object is in a correct format
SYSTOOL	JSON2BSON	Convert JSON strings into a BSON document format
SYSTOOL	JSON_GET_POS_ARR_INDEX	Find a value within an array
SYSTOOL	JSON_LEN	Returns the count of elements in an array type inside a document
SYSTOOL	JSON_TABLE	Returns a table of values for an array field
SYSTOOL	JSON_TYPE	Returns the data type of a specific field within a JSON document
SYSTOOL	JSON_UPDATE	Update a field or document using set syntax
SYSIBM	JSON_VAL	Extracts data from a JSON document into SQL data types

# Technical Preview: Indexes on Column Organized Tables

## ▪ Currently Db2 BLU has limited index support

- Unique indexes *implicitly created* for PRIMARY/UNIQUE KEY constraints
- A unique index can be used to access data if *at most 1 row qualifies*
  - Every key in the index must have an equality predicate OR
  - FETCH FIRST 1 ROW ONLY clause is specified

## ▪ Indexes improve performance for a wide range of applications:

- Transactional (OLTP): The “**T**” in HTAP
- **Operational analytics**: ETL/ELT with transactional characteristics
- Analytics: Some ‘heavy lifting’ queries with certain access patterns

## ▪ Indexes on BLU when delivered (TBD) is initially targeted for operational analytics

- Especially SAP BW ETL
- Future HTAP work will provide further performance improvements

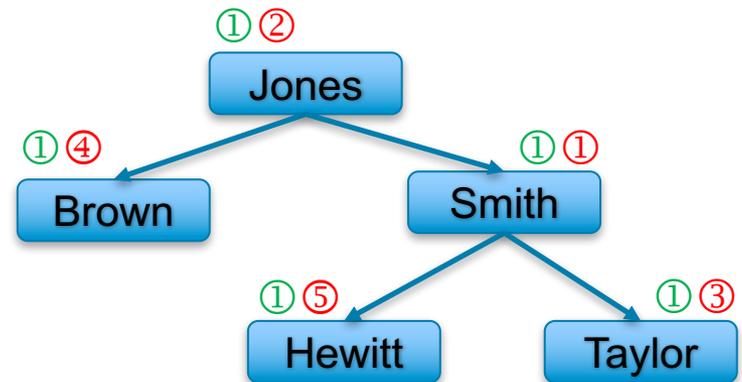
# Technical Preview: Indexes on Column Organized Tables

## ▪ In Db2 11.1.2.2, Indexes are provided as a technical preview

- Non-production use only – **not supported**
- Requires 2 registry variables to be set
  1. **DB2\_BLU\_OLTP=INDX:Y**
    - Allows **CREATE** and **DROP INDEX**
  2. **DB2\_EXTENDED\_OPTIMIZATION="COL\_ISCAN IXONLY\_UR"**
    - Enables use of indexes for **UR isolation, index only access queries only**

## ▪ Limitations in technical preview:

- CREATE INDEX and REORG INDEX .. REBUILD mode will allow concurrent readers but not concurrent writers.
- Jump scan
- Deferred fetch index plans (index ANDing, ORing and list prefetch)
- Star join and zigzag join
- Index extensions (most commonly used by spatial extenders)
- Expression-based indexes
- Scan sharing
- Intra-partition parallel scans\*\*
- RANDOM key order
- EXCLUDE NULL KEY option
- INCLUDE columns\*\*
- WITH SPECIFICATION ONLY clause (applies to nicknames)
- CLUSTER (incompatible with BLU implementation)

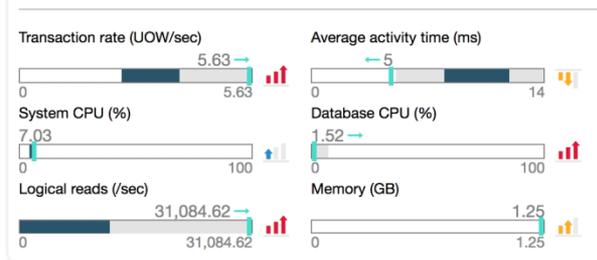


## ▪ More details will be in a tech note

## Monitoring and Alerts

Accelerate hybrid problem determination

- Redesigned home page card view
- Active **hotspot drill down** for in depth monitoring
- dashDB Appliance integration



## Security

Deliver mainframe class security for DSM

- Data Server Manager Login authentication through **Kerberos**
- **Multi factor** authentication support for Data Server Manager
- Publish end to end security guide



I know



I am



I have

## Development and data management

Simplify application development and data management

- Hybrid logical warehouse: make remote data **look and act like local tables**
- Add new data sources for **Fluid Query**
  - **Apache Hive**
  - **Cloudera Impala**
  - IBM BigInsights
  - IBM Db2 and IBM dashDB
  - IBM Db2 for z/OS
  - **IBM PureData System for Analytics (Netezza)**
  - **Microsoft SQL Server**
  - Oracle (using ODBC client)
  - **Oracle (using NET8 client)**
  - **Teradata**
  - **JDBC (generic data sources)**



## New Tuning Enhancements

- Several usability enhancements in the tuning wizard
- Manage filters for capturing data for tuning
- Compare Access Plan to show differences between original and new access plans
- Improved Index Advisor Tool

STMTNO	STMT_TYPE	STMT_TEXT	TOTAL_COST	PROCCS	PROCCS	COLLID
1	DECLARE CURSOR	DECLARE CURSOR CURSOR CURSOR	-1	0	0	0
2	DECLARE CURSOR	DECLARE CURSOR CURSOR CURSOR	-1	0	0	0
3	DECLARE CURSOR	DECLARE CURSOR CURSOR CURSOR	-1	0	0	0
4	DECLARE CURSOR	DECLARE CURSOR CURSOR CURSOR	-1	0	0	0
5	DECLARE CURSOR	DECLARE CURSOR CURSOR CURSOR	-1	0	0	0
6	DECLARE CURSOR	DECLARE CURSOR CURSOR CURSOR	-1	0	0	0
7	DECLARE CURSOR	DECLARE CURSOR CURSOR CURSOR	-1	0	0	0
8	DECLARE CURSOR	DECLARE CURSOR CURSOR CURSOR	-1	0	0	0
9	DECLARE CURSOR	DECLARE CURSOR CURSOR CURSOR	-1	0	0	0
10	DECLARE CURSOR	DECLARE CURSOR CURSOR CURSOR	-1	0	0	0

# Db2 Connect, DS Driver and DS Gateway 11.1.2.2

## Expanding Cloud Environments

- **.Net Core and EF** core support expands .Net workloads to private cloud also.
- Driver serviceability simplification



## Performance and Administration

- **Speed application access** to CAE envs
  - **External Tables** expands access to large data
  - Requires applicable Data Store support . Ex. delivered in June dashDB Local update
- DAA class improvements (Internal performance improvements)



## Expanding IBM Environments

- **IMS T4 Connectivity** enables communication from IMS to distributed applications
- **IBM I OS V7R3** support enables access to Db2 on i



## Developers

- **VisualStudio 2017 Beta Q3** promotes developer usage of DS Driver and Db2 with Common Analytics Engine



# Advanced Recovery Feature for Db2 11.1.2.2

## Further Db2 Encryption Support

- Support for Db2 native encryption having their keystore stored on a KMIP server
- Exploit encryption features that are built right into the database with Db2 Advanced Recovery Feature
- Unload in native mode from an encrypted database, if the table is located into a tablespace having several containers



## Improved BLU Support for HPU

- Improve the support of column organized tables when unloading from Db2 11.1.1.1 and above



## Enhance Db2 Merge Backup for Log Processing

- Enhanced so that when restoring a Db2 Merged Backup containing logs
- Maximize use case of Db2 Merge Backup and TSM fixing problems with bad Db2 log files
- Fixed problems with merging compressed backups



## More Enhancements and Features

- HPU: Unload from a backup when using an older version of Db2 than the one of the Db2 instance concerned



# Agenda

- Updates to the Db2 Brand
- 11.1.2.2 Licensing and Packaging
- Download and Go
- 11.1.2.2 Content Details
- **Roadmap Overview and Additional Resources**

# Db2 Roadmap



**Leading In-Memory HTAP**

**Deeper BLU Optimizations for Operational Workloads**

- Performance enhancements
- Builds on 4Q '16 advances
- Enables use of BLU beyond strictly analytic workloads

**BLU for Operational Analytics**

- Secondary indexes enabling large scale PDOA and ODS workloads

**BLU for OLTP**

- Added performance and function
- More index exploitation and function (e.g. cursored operations)
- Enables use of BLU in straight-forward OLTP applications

**BLU for advanced OLTP**

- Added performance and function, including triggers
- Enables use of BLU in advanced OLTP workloads

**Mainframe Class Availability, Robustness, Security**

**Near-zero outage recovery**

- Online crash recovery
- pureScale REBUILD restore

**Zero outage maintenance**

- No outage for key maintenance operations including:
- Space reclaim for pureScale
- On-line INDEX create for pureScale

**Dramatic reduction in backup and recovery windows**

- Crash recovery speed improvements
- Proactive tablespace growth
- Back-up manageability/speed enhancements
- Security advances: e.g. more tolerance of key store manager

**All topology changes fully online**

- Full online topology changes for pureScale
- Improved online-schema change
- Security: KMIP 1.2 Support
- Schema-level security

**Future Proofed**

**Native JSON support**

- JSON SQL support

**REST API**

- REST support
- Enhanced JSON/SQL

**Spark integration**

**JSON SQL standard support**

# Sign up for early access: [ibm.biz/DB2-EAP](http://ibm.biz/DB2-EAP)

[Marketplace](#)
Search

IBM Analytics > Technology > Data and content management > DB2 > Early Access Program >

## DB2 Early Access Program

↓
Download Now

The Early Access Program offers a sneak peek at next generation IBM DB2 technology. Take this opportunity to shape the future of DB2, test integration with your existing infrastructure, and be part of the community that makes DB2 great.

### Shape the future of DB2

Provide feedback and shape the world's most advanced database software. We are constantly improving and innovating upon DB2 and your feedback is critical to our success. Your opinion matters to us.

### Be Part of a Community

Join our community of business leaders, developers, and enthusiasts that is the DB2 Early Access Program community. Discuss ideas, share knowledge, and help each other in all things DB2 related.

### Gain Early Access

Be one of the first customers to use next generation DB2 software. Whether you are upgrading from DB2 10.x or 9.x, or have never used DB2 before, this is something for everyone. Gain a glimpse into the future of database software.

# Stay connected with our blogs

ibmdatamanagement.co

**IBM Data Management Blog**  
The latest news and information from the IBM Data Management Team

HOME | IBM DB2 | DB2 BLU ACCELERATION | DB2 IN THE NEWS | BIG DATA & ANALYTICS | UNCATEGORIZED | EVENTS | INFORMIX | IOT | DB2 TOOLS

## Why it is a great time to migrate to DB2 11.1?

MARCH 6, 2017 | LEAVE A COMMENT

by Sajjan Kuttappa, Content Marketing Manager, IBM DB2

IBM released the latest version of its DB2 software – DB2 11.1 for Linux, UNIX and Windows in 2016. Improvements were introduced including simplification of the upgrade process. Recognizing that many have databases on older releases, DB2 11.1 increased the number of back-level releases that can be included in Version 9.7. This means if you are on DB2 9.7, you would not have to go through another version or version 10.5, to upgrade to DB2 11.1. This allows you to start benefiting from the capabilities in DB2 Windows version 11.1, with significantly less time and effort.

For customers who have DB2 version 9.7 or version 10.1, this is a great time to upgrade to the new releases as they are due to go out of support on the 30<sup>th</sup> of September, 2017. For more details visit [this link](#).

Apart from the upgrade process itself, DB2 11.1 introduced several improvements that allow DB2 to handle larger workloads. Additional advances were introduced in BLU Acceleration technology including the ability to handle several hundred servers. This feature extends BLU's in-memory processing speeds to a massive Petabyte of cores. DB2 11.1 also offers other core benefits like improved security mechanisms, improved availability and performance of your queries. Check out DB2 11.1 highlights [here](#).

Learn more about how you can benefit by migrating to DB2 11.1 by visiting our [website](#).

ibmbigdatahub.com

**IBM Big Data & Analytics Hub**

ABOUT | CONTACT | CONTENT BY TYPE | SUBSCRIBE

EXPLORE BY TOPIC: USE CASES | INDUSTRIES | ANALYTICS | TECHNOLOGY

EVENTS | FOR DEVELOPERS | BIG DATA & ANALYTICS HEROES

## Go beyond collaboration for data-driven success

READ THE BLOG

### The Latest

- ANALYTICS**  
VIDEO | 2:08  
The smarter payment platform for immediate payments
- BIG DATA TECHNOLOGY**  
BLOG | by Andrea Braida  
Fundamentals for sure-fire cloud data warehouse optimization
- ANALYTICS**  
WHITE PAPERS & REPORTS  
2017 Gartner Magic Quadrant for Data Science Platforms

### Most Popular

- BLOG**  
Four perspectives on data lakes
- BLOG**  
How to leverage the power of prescriptive analytics to maximize the ROI
- BLOG**  
Wanted: Smart humans
- BLOG**  
Cloud Data Center



# Get involved with a local Db2 community through IDUG

## Experience IDUG®

*The Worldwide DB2 User Community*



### Educational Events

Quality Education

Volunteer committees of DB2 professionals plan the IDUG educational events

Networking Opportunities with Top Product Developers, Industry Experts, Renowned Consultants, Fellow Users, and Vendors

### IDUG Mentor Program

Opportunity for IDUG members to pass on valuable skills with an 80% discount to bring a coworker to an IDUG DB2 Tech Conference

Find a local user group chapter!

<http://bit.ly/1z2RYcl>



Join IDUG Now!



[www.IDUG.org/join](http://www.IDUG.org/join)

**IBM**®

## Legal Disclaimer

- © IBM Corporation 2017. All Rights Reserved.
- The information contained in this publication is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this publication, it is provided AS IS without warranty of any kind, express or implied. In addition, this information is based on IBM's current product plans and strategy, which are subject to change by IBM without notice. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this publication or any other materials. Nothing contained in this publication is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.
- References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates. Product release dates and/or capabilities referenced in this presentation may change at any time at IBM's sole discretion based on market opportunities or other factors, and are not intended to be a commitment to future product or feature availability in any way. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.
- If the text contains performance statistics or references to benchmarks, insert the following language; otherwise delete:  
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.
- If the text includes any customer examples, please confirm we have prior written approval from such customer and insert the following language; otherwise delete:  
All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.
- Please review text for proper trademark attribution of IBM products. At first use, each product name must be the full name and include appropriate trademark symbols (e.g., IBM Lotus® Sametime® Unyte™). Subsequent references can drop "IBM" but should include the proper branding (e.g., Lotus Sametime Gateway, or WebSphere Application Server). Please refer to <http://www.ibm.com/legal/copytrade.shtml> for guidance on which trademarks require the ® or ™ symbol. Do not use abbreviations for IBM product names in your presentation. All product names must be used as adjectives rather than nouns. Please list all of the trademarks that you use in your presentation as follows; delete any not included in your presentation. IBM, the IBM logo, Lotus, Lotus Notes, Notes, Domino, Quickr, Sametime, WebSphere, UC2, PartnerWorld and Lotusphere are trademarks of International Business Machines Corporation in the United States, other countries, or both. Unyte is a trademark of WebDialogs, Inc., in the United States, other countries, or both.
- If you reference Adobe® in the text, please mark the first use and include the following; otherwise delete:  
Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
- If you reference Java™ in the text, please mark the first use and include the following; otherwise delete:  
Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.
- If you reference Microsoft® and/or Windows® in the text, please mark the first use and include the following, as applicable; otherwise delete:  
Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.
- If you reference Intel® and/or any of the following Intel products in the text, please mark the first use and include those that you use as follows; otherwise delete:  
Intel, Intel Centrino, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
- If you reference UNIX® in the text, please mark the first use and include the following; otherwise delete:  
UNIX is a registered trademark of The Open Group in the United States and other countries.
- If you reference Linux® in your presentation, please mark the first use and include the following; otherwise delete:  
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both. Other company, product, or service names may be trademarks or service marks of others.
- If the text/graphics include screenshots, no actual IBM employee names may be used (even your own), if your screenshots include fictitious company names (e.g., Renovations, Zeta Bank, Acme) please update and insert the following; otherwise delete: All references to [insert fictitious company name] refer to a fictitious company and are used for illustration purposes only.