# DB2 V11.1.1.1 Overview



© 2017 IBM Corporation

## **Safe Harbor Statement**

Copyright © 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

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES 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, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON CURRENT THINKING REGARDING TRENDS AND DIRECTIONS, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. FUNCTION DESCRIBED HEREIN MY NEVER BE DELIVERED BY I BM. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. 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.

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 ™), 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

## Agenda

- What is DB2 11.1.1.1?
- What's new in DB2 11.1.1.1
  - Encryption Enhancements
  - Even Greater Availability
  - Additional Core Functionality
  - Even Greater SQL Compatibility
  - Increased Serviceability
  - Performance Improvements
  - Removal of Limitations
  - Support for Additional Operating Systems
  - Updates to Data Server Manager
  - Updates to the Data Server (DS) Driver and DB2 Connect
- What to know before updating from DB2 11.1 GA



http://ibm.box.com/v/DB2v11eBook

## Why 11.1.1.1 (instead of 11.1.1)?

- The official DB2 product signature consists of 4 parts and has the format VV.RR.MM.FF where:
  - VV = Version number
  - RR= Release number
  - MM = Modification number
  - FF = Fix pack number
- Until now, the modification value for DB2 LUW has always been 0 (zero)
- Traditionally, interfaces that return the product signature have supplied only 3 elements – VV, RR, and FF
- It has not always been obvious when a Fix Pack contains new functionality
- Starting with this deliverable, the modification value (MM) will be used to indicate the existence of new functionality or behaviors in a fix pack

#### Introduce 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



## **DB2 Developer-C Supporting Programs**

#### Sample of Supporting Programs to be included with 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
- <sup>9</sup> IBM WebSphere MQ

## DB2 11.1.1.1 Highlights

#### **Comprehensive Enterprise Security**

#### **Enterprise Encryption**

- PKCS#11 HSM support
- SSL Encryption for HADR
  - Initially available on Linux/x86 only

#### **Higher Availability and Core Capabilities**

#### **Even Greater Availability**

- Seamless HADR pureScale upgrades
- FORCE\_ALL for ADMIN\_MOVE\_TABLE() Asynchronous UNDO (technical preview)

#### **Additional Core Functionality**

- Workload Manager (WLM) enhancements
- Federation integration and simplification
- DECIMAL and DECFLOAT improvements

#### More Compatibility and Serviceability

#### Even Greater SQL Compatibility

- Column support for **BOOLEAN** data type
- Common table expressions (DB2 for z/OS)
- WITH and SELECT INTO support

#### **Increased Serviceability**

- Maximum log file size increased to 64 GB
- Client info in db2diag for lock and lock errors
- Backup image size in event monitor history

#### **Column-Organized (BLU) Tables**

#### Performance Improvements

- Synopsis table enhancements (HTAP)
- Additional SIMD exploitation
- INSERT from sub-select performance improvements
- Aggregation enhancements, sort elimination
- Additional advantages for SAP

#### **Added Function – Removing Limitations**

- Automatic Dictionary Creation on uncommitted data
- ALTER VARCHAR/VARGRAPHIC length support



#### Additional Operating System Support



#### DB2 pureScale

• SLES 12, RHEL 6.8 on x86

#### Ubuntu on z Systems

- Ubuntu 16.04 on zLinux (non-pureScale)
- Text Search support

#### Windows

Support for Windows Server 2016

## **DB2 Native Encryption HSM Support**

- Support for PKCS #11-compliant Hardware Security Modules (HSMs) is now available for:
  - Gemalto (formerly Luna) Safenet HSM (firmware version 6.23.0; software version 6.1 and above)
  - Thales nShield Connect+ (software version 11.50)



## **SSL Encryption Between HADR Primary and Standby Servers**

- Provides integrated protection of sensitive data in the log stream
- Enabled via the HADR\_SSL\_LABEL database configuration parameter
- All synchronization modes supported
- Multiple standbys supported
- Currently works with Linux on x86 (nonpureScale) platforms only



## **Seamless HADR Upgrades for DB2 pureScale Environments**

- DB2 pureScale HADR environments can now be upgraded without having to stop HADR or re-initialize the standby(s)
- Standby databases 'replay' upgrades made to the primary
- Not supported with DB2 9.7, 10.1, or 10.5 Fix Pack 8 or earlier; will be supported in an upcoming 10.5 Fix Pack



## FORCE\_ALL Option For ADMIN\_MOVE\_TABLE()

- The ADMIN\_MOVE\_TABLE() procedure provides the ability to move data from an active 'source' table to a new table object that has the same name, while the table remains online and accessible.
- A new FORCE\_ALL option enables administrators to avoid potential lock issues when performing table movement operations:

```
ADMIN_MOVE_TABLE(<SchemaName>,<TableName>,'','FORCE_ALL','SWAP`)
```

- Helps avoid deadlocks and timeouts due to heavy use of the source table during the SWAP phase
- Behaves like the FORCE APPLICATION command
- Requires SYSADM, SYSCTRL, or SYSMAINT authority
- Care should be exercised when using the FORCE\_ALL option!

## WLM Dispatcher: CPU Control at the Database Level

- New database-level share values control the division of CPU resources between databases
  - Enables prioritized allocation of CPU in multi-database environments
  - Configured via the WLM\_CPU\_SHARES and WLM\_CPU\_SHARE\_MODE database configuration parameters
- New database-level limit value controls the total CPU consumption of all service classes within a database
  - Configured via the WLM\_CPU\_LIMIT database configuration parameter



## **Federation Integration and Enhancements**

### Federation capabilities extended to more data sources

- More drivers directly integrated into the DB2 installer
- Expansion to both SQL- and NoSQL-based data stores
- Expansion to both on-premises and cloud data sources

### Set-up simplified

- CREATE WRAPPER statement and WRAPPER clause of CREATE SERVER now optional (a default wrapper will be created if clause is omitted)
- Can create server with host, port, and service name without first cataloging the remote node and database



© 2017 IBM Corporation

## Core Functionality and SQL Compatibility Improvements<sup>(1)</sup>

### DECIMAL/DECFLOAT data type improvements

- Arithmetic on AIX, Intel, Linux on Power (LE), and zLinux accelerated
- AIX and Intel Conversion between row- and column-organized formats optimized
- SUM() operations on DB2 for zLinux optimized

### Table column and expression support for BOOLEAN data type added

- Valid values include TRUE, 'true', 't', 'yes', 'y', 'on', '1', FALSE, 'false', 'f', 'no', 'n', and '0'
- Some initial restrictions:
  - Replication (SQL / Q Replication, CDC) not allowed
  - IMPORT operations not allowed
  - ALTER NICKNAME column data type change to/from BOOLEAN not allowed (column name changes supported)
  - Cannot be used with ANALYZE\_TABLE expressions
  - External functions and procedures written in C, COBOL, CLR, and OLE cannot use BOOLEAN input or return values

## Core Functionality and SQL Compatibility Improvements<sup>(2)</sup>

- Common Table Expression support extended to increase compatibility with DB2 on z
  - The WITH clause can now be used with SELECT INTO statements for Static SQL, SQL User-Defined functions, and SQL Stored Procedures:

```
WITH cte_1(c1, c2) AS (SELECT c1, c2 FROM t1)
SELECT SUM(c2) INTO :outvar FROM cte_1
WHERE c1 BETWEEN :lowval AND :highval
```

- The db2\_install command now requires the user to accept the license agreement before the installation can proceed
  - Acceptance can be made by specifying the new -y parameter
  - If the -y parameter is not specified, the user is prompted to either accept or decline the license agreement
  - The -y parameter must be specified if the -n parameter is specified

### Increased Serviceability<sup>(1)</sup>

### Log file size limit increased to 64 GB (from 4 GB)

- Larger log files allow for larger log space
- Actual size used is controlled by the LOGFILSIZ database configuration parameter
- Client information is now written to db2diag.log for lock and lock error events
  - Additional CLIENT diagnostic information (i.e., CLIENT USERID, CLIENT ACCTNG, CLIENT APPLNAME, and CLIENT WRKSTNNAME) is recorded in the log when:
    - A transaction exceeds the number of log files it can span (NUM\_LOG\_SPAN database configuration parameter value)
    - A transaction exceeds the amount of primary log space it can consume (MAX\_LOG database configuration parameter value)
    - A transaction triggers a lock escalation event
  - In MPP (DPF) environments, this information is written only if the events occur on the coordinator node

### Increased Serviceability<sup>(2)</sup>

- Backup image sizes can now be obtained using an event monitor
  - New BYTES\_TOTAL monitor element added to the UTILSTOP logical data group
  - Used by BACKUP change history event monitors only

### Example:

 Connect to a database named TEST (that has been configured for online backups), create a BACKUP change history event monitor, and then start an online backup:

```
$ db2 "CONNECT TO test"
$ db2 "CREATE EVENT MONITOR em1 FOR CHANGE HISTORY WHERE EVENT IN
(BACKUP) WRITE TO TABLE AUTOSTART"
$ db2 "BACKUP DATABASE test ONLINE"
```

- Obtain the size of the backup image created using SQL:

\$ db2 `SELECT UTILITY_TYPE, START_EVENT_TIMESTAMP, BYTES_TOTAL FROM utilstop_em1"					
UTILITY_TYPE	START_EVENT_TIMESTAMP	BYTES_TOTAL			
ВАСКИР	2017-01-10.8.43.48.663294	184627200			

© 2017 IBM Corporation

## Synopsis Updates with Block Inserts (Good)

	First	Last	Address	City	State
First Last Address City State					
John Smith 1 Cork St Carson CA (2)					
Katrina       Jones       132 Front Ave       Buffalo       NY       3         (4)					
Bob   Taylor   23 Bee St   Akron   OH     5					
Enzo Brown 9 Elbow Pl Denver CO					
Rick Hewitt 31 King St Astoria OR					
	First	Last	Address	City	State
	Bob Rick	Brown Taylor	1 Cork St 9 Elbow Pl	Akron Denver	CAOR

- 1) Records are inserted into the various columns
- 2) The synopsis table is updated with MIN/MAX values for each column
- 3) Only one synopsis update for the block insert

## Synopsis Updates with Individual Inserts (Not so good)

	First	Last	Address	City	State
First Last Address City State					
John Smith 1 Cork St Carson CA (2)					
Katrina     Jones     132 Front Ave     Buffalo     NY       (4)					
Bob   Taylor   23 Bee St   Akron   OH     5					
Enzo Brown 9 Elbow Pl Denver CO					
Rick Hewitt 31 King St Astoria OR					
	First	Last	Address	City	State
	John John John Katrina	Smith Smith Jones Smith	1 Cork St 1 Cork St 1 Cork St 132 Front Av	Carson Carson Buffalo Carson	CA CA CA NY
	Bob Katrina Bob Katrina	Jones Taylor Brown Taylor	1 Cork St 23 Bee St 1 Cork St 9 Elbow Pl	Akron Carson Akron Denver	CA OH CA OH
	Bob Rick	Brown Taylor	1 Cork St 9 Elbow Pl	Akron Denver	CAOR

- 1) Records are inserted into the various columns
- 2) The synopsis table is updated with MIN/MAX values for each column
- 3) One synopsis update for the each insert

## **New In-Memory Synopsis**

- Synopsis updates collected in memory until 1024 rows/TSNs have accumulated
- The transaction that inserted 1024th value will update the persisted synopsis table with a single entry
  - After recovery from crash, 'lost' synopsis values rebuilt via limited scan of 'tail' of table
- New in-memory algorithm used for all synopsis tables in DB2 v11.1.1.1 and beyond
- Benefits:
  - Significant reduction in pages updated in small transactions and corresponding significant performance improvements
  - Significant storage savings in synopsis tables for small transactions
  - Significant performance improvements for queries that exploit synopsis tables

# **In-Memory Synopsis : Example Benefits**

- Previous synopsis maintenance was designed for batch inserts (1000s of rows)
- In-memory synopsis is aimed at <u>both</u> batch inserts and small OLTP transactions
  - Updates values in memory and writes them to synopsis table after 1024 rows are covered

#### In this lab test we observed Elapsed Time of INSERT Size of Synopsis Table **Query Elapsed Time** With and Without In-With and Without In-With and Without In-Memory Memory Synopsis Memory Synopsis **Synopsis** 600,00 900 0.07000 pages) Time (seconds) Time (seconds) 800 0.06000 500.00 700 0.05000 400,00 豊 600 46% 60% 0.04000 ble 500 300.00 0.03000 Та 400 00,000 Elapsed Elapsed <sup>-</sup> 0.02000 Synopsis 300 200 0.01000 100 0,0000 Onery Insert 0,00 EE of MPP EE MPP 0 Size With No Optimization No Optimization With Optimization No Optimization With Optimization Optimization

- IBM p760 / POWER7+ 32 cores / 1TB RAM
- Table with 50 columns.

- Single user
- Insert test 100K rows with commit count = 1
- Select query with 10 pairs of range predicates

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will 24 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 and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.



## **Multi-Core Parallelism Directions**

- DB2 has exceptional multi-core scalability for queries:
  - Very strong scalability on the largest SMPs
  - Combined MPP and SMP parallelism
  - Includes the SELECT component of an INSERT from sub-select
  - Includes columnar and row-based tables

#### Future Deep Multi-Core Parallelism for Writes

- Includes a variety of data population methods (i.e., INSERT, INGEST, etc.)
- Includes index update parallelism
- Includes MPP support

### • DB2 v11.1.1.1 adds INSERT parallelism

- Very significant reduction in data population/ingest job duration
- Very significant reduction in ETL/ELT batch jobs
- UPDATE/DELETE expected in the future



## **Multi-Core INSERT Parallelism – Example Benefits**

## Elapsed time (seconds) for:

INSERT INTO table2 (SELECT \* FROM table1)



Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will 26 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 2019 BMage or poration configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

## **Multi-Core INSERT Parallelism – Details**

Enabled with two registry variables – DB2\_REDUCED\_OPTIMIZATION and DB2\_EXTENDED\_OPTIMIZATION:

\$ db2set DB2\_REDUCED\_OPTIMIZATION=ENABLE\_RTABLE\_INS -immediate

\$ db2set DB2\_EXTENDED\_OPTIMIZATION=CDE\_PAR\_IUD -immediate

#### Remember to keep previous settings, if needed; for example:

\$ db2set DB2\_REDUCED\_OPTIMIZATION=STARJN\_CARD,ENABLE\_RTABLE\_INS -immediate

#### In the future, we expect this behavior to be enabled by default

- Once enabled, the degree of parallelism is automatically determined based on data volume (i.e., number of rows)
  - INSERT with values that touch just a few rows is not likely to be parallelized
  - INSERT with sub-select from a large source table is likely to be parallelized
  - The actual DEGREE of parallelism chosen by the optimizer is shown in the **RETURN** operator of the **EXPLAIN\_ARGUMENT** table

For more information, refer to the Technote titled **Parallel INSERT** at http://www-01.ibm.com/support/docview.wss?uid=swg27049357.

## Other BLU Enhancements<sup>(1)</sup>

### Additional SIMD enhancements

- Exploitation of the latest SIMD technology:
  - Intel® Advanced Vector Extensions 2 (Intel® AVX2)
  - Streaming SIMD Extensions 2 (SSE2)
  - Supplemental Streaming SIMD Extensions 3 (SSSE3 or SSE3S)
  - Streaming SIMD Extensions 4 (SSE4)
- Support extended to Power8, z13<sup>™</sup> and Intel CPUs supporting AVX2
- More functions and bit operations
- Utilization of the Fletcher-64 checksum algorithm



## Other BLU Enhancements<sup>(2)</sup>

- Performance improvements for INSERT from sub-select operations against BLU tables
  - With IDENTITY columns up to 140x faster
  - In MPP environments where RANDOM data distribution is used up to 3.5x faster
- Reduction (or elimination) of sort operations in some access plans (e.g., queries against BLU tables with OLAP specifications and/or ORDER BY clauses)
- Use of uncommitted data for Automatic Dictionary Creation for BLU tables
  - Compression dictionaries are made available earlier for:
    - INSERT, IMPORT, and INGEST operations
    - Large INSERT operations such as INSERT with sub-selects
  - Higher PCTENCODED (percent encoded) values in SYSCAT.COLUMNS results in:
    - Better query performance
    - Increased throughput for high concurrency workloads

## Other BLU Enhancements<sup>(3)</sup>

- Aggregation\* enhancements for improved memory usage and performance:
  - Improved decisions to pre-partition using the GROUP BY columns of the aggregation input stream
  - Better memory and performance for single/unique group results
    - If a group contains a single row, then the column value itself is returned as the aggregation result
  - Better decisions to avoid a "Partial-Final" MPP GROUP BY when the partial GROUP BY does not significantly reduce the aggregation stream
  - New and improved 64-bit hashing algorithm reduces the number of collisions in the hash table and better utilizes memory needed for aggregation

#### \* Data aggregation is any process in which information is gathered and expressed in a summary form.

- Examples of DB2 11.1.1.1 benchmarking (unofficial) versus DB2 11.1 GA
  - TPC-DS shows a 24% improvement
  - TPC-H shows a 14% improvement
  - SAP-BW shows a 22% improvement



While NOT part of any formal and official benchmarking efforts, the measurements and results on this page show the performance benefits (elapsed time) of some workloads that have been tested internally in the lab. There is no way you are ever going to get the same results on your systems so don't hold us to it.

IEM 🕷

## **Removal of Restrictions and Limitations**

- A system temporary table space having a page size of 32 K is no longer required for extended row sizes
- The ability to alter the length of VARCHAR/VARGRAPH columns in column-organized tables is now supported
  - String units (e.g. CODEUNIT32) cannot be altered
- The db2convert command can now be used to convert roworganized tables in MPP and with generated identity columns to column-organized tables
- System maintained materialized query tables (MQTs) in MPP can now reference a nickname



- DB2 pureScale support for:
  - SLES 12, SP1 on x86
  - RHEL 6.8 on x86
- Non-pureScale support for Windows Server 2016
  - No warehouse components or Tivoli System Automation
- Non-pureScale support for Ubuntu 16.04 on z Systems
  - 11.1 GA had prerequisites on 3 libraries which had to be installed manually; this was resolved in 11.1.1.1
- Text Search support added for all Linux on z Systems





iem 🔉

## **New DB2 JSON Functions**

#### Several JSON functions now available on DB2 10.5 and DB2 11.1

- Insert, Retrieve, and Document Validation
  - BSON2JSON Convert BSON formatted document into JSON strings
  - JSON2BSON Convert JSON strings into a BSON document format
  - BSON\_VALIDATE Checks to make sure that a BSON field in a BLOB object is in a correct format
- Field Retrieval
  - JSON\_VAL Extracts data from a JSON document into SQL data types
  - JSON\_TYPE Returns the data type of a specific field within a JSON document
- Array Retrieval
  - JSON\_TABLE Returns a table of values for a document that has array types in it
  - JSON\_LEN Returns the count of elements in an array type inside a document
  - JSON\_GET\_POS\_ARR\_INDEX Find a value within an array
- Document Update
  - JSON\_UPDATE Update a field or document using set syntax

## **IBM Data Server Gateway for OData Version 1.0.0**

### • What is the problem?

- Customers want to access Data from non traditional platforms like Mobile platforms, from mobile apps (Android, Windows, iOS, etc.)
  - And in some cases want to publish access to data as well to 3rd parties
- Mobile and Cloud developers expect direct HTTP and JSON (API based) access to data (internal like DB2, 3rd party like Salesforce, SAP) – without requiring Database Drivers

#### • What is being delivered:

- ODATA is standards based way for our customers to leverage DB2, and dashDB in modern Cloud and Mobile Application Models
- Delivery mechanism to be synergistic with DS Driver packages with DB2 V11 Mod 1
- https://www.youtube.com/watch?v=z06bl\_K6Ckc
- https://www.youtube.com/watch?v=dHVBiTphoCM

## Preview: Online Crash Recovery (aka Async UNDO)

- Significantly increase database availability during crash recovery
  - Focus on scenarios where large batch operations are updating the database
- Allow database connections during the UNDO phase of crash recovery (Also HADR TAKEOVER BY FORCE)
  - The UNDO phase is significantly longer if a crash occurred when batch operations or other long-running (e.g. errant) transactions were executing

### Alter crash recovery to:

- Acquire exclusive locks on all tables or table partitions with in-flight operations immediately after the REDO phase
- Allow new connections to the database as early as possible in the UNDO phase so:
  - Access to unlocked tables/partitions can proceed normally
  - UR access to locked tables/partitions is allowed
- Process the UNDO phase asynchronously, releasing table locks as UNDO progresses

IBM Analytics

#### **Technical Preview: Online Crash Recovery (Async UNDO)** DB up pS No Change prior to DB up DB down 11.1 11.1.1.1 <u>DB down</u> <u>DB up (\*)</u> DB up **Preview** UNDO - REDO -45000 40000 35000 pureScale .... 30000 **Fotal Transactions Per Second** 25000 **Preview** Current 20000 15000 **Connections allowed after synchronous** portion of UNDO completes. Length of 10000 synchronous portion is dependent on actual data operations. More detail 5000 In tech note referred to on next page. <sup>37</sup> (\*) Some data may be locked while UNDO proceeds concurrently © 2017 IBM Corporation Time

## **Preview: Online Crash Recovery – Details**

#### To enable the Online Crash Recovery Technical Preview:

 Assign the value 'YES' to the undocumented registry variable DB2\_ASYNC\_UNDO:

db2set DB2\_ASYNC\_UNDO=YES

- Reactivate the database IMPORTANT: Use in DB2 11.1.1.1 non-production environments only; lab validation is not complete!
- 3. Once enabled, the following administrative notification message is written during crash recovery when the aysnchronous portion of the UNDO phase begins and database connections are allowed:

#### ADM1505I Crash recovery has completed synchronous processing

For more information, refer to the Technote titled **Database accessibility during Backward phase of crash-recovery or 'HADR Takeover by Force'** at: http://www-01.ibm.com/support/docview.wss?uid=swg21994342

## Data Server Manager 2.1.2 – What's New for DB2 11.1.1.1

#### Manage your IBM data Enterprise

- Enterprise view of DB2 and dashDB
   KPIs and alerts from one dashboard
- Data Server Manager Login authentication through LDAP
- New enterprise grid view shows hundreds of DB2 and dashDB servers at a glance

Name 🔹 🚺 Refresh Ev	very 6	0s 🔹	Ç				Q	Ծ・⊞・ℰ・目
Name		Alerts	•	SYS CPU	VO (/sec)	Memory (GB)	Transaction Rate (UOW/sec)	Storage Accesses (hour)
History Repository	<b>~</b>			14.33%	313.53 🕇	0.36	7.02	260.30K
dashDB Local	<b>*</b>	1 🛃		33.36%	545.50	3.17	78.21	631.98K
dashDB Transaction	<b>~</b>			48.14%	4057.94	1.60	83.07 🕇	2.07M

### **Ensure DB2 High Availability**

- See the availability of all of your HADR and DB2 pureScale clusters at a glance
- Dashboard, SNMP and Email alerts that highlight potential availability problems before they happen



### **Monitor In Depth**

- New powerful database performance overview page for rapid problem determination
- Database time breakdown
- 18 KPIs in one screen for real-time or historical analysis
- Breakdown by workload
- For DB2 and dashDB



#### Making remote data act like local data

- Make data in remote DB2, BigSQL or dashDB databases look and act like local tables with 11.1.1.1 and Fluid Query
- Explain and monitor queries with calls to remote data sources with 11.1.1.1



#### **International DB2 Users Group**



## IDUG DB2 Tech Conference Lisbon, Portugal October 1-5, 2017

- The premier European DB2 Conference will take place on October 1-5, 2017 at Lisbon, Portugal. Make your plans to attend and experience the latest in DB2 technologies, networking opportunities and the technical content you need to be successful.
- The 2017 event offers new educational opportunities and more training than ever before! Attendees will experience: 5 days of education sessions
- Half and full-day workshops
- More than 100 one-hour technical sessions
- Three expert panels on z/OS, LUW & Application Development
- And much more!
- Register at <u>www.idug.org</u> before 7 August 2017 to receive *Early Bird Discount!*
- Conference hotel: EPIC SANA Lisboa Hotel, Lisbon

