



**CenCan /
CCDUG /
DTS**

Db2: Latest from the Lab

Keri Romanufa, IBM

L01: Monday Sept 19, 10:15

Toronto, ON

Please Note:

- IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice and at IBM's sole discretion.
- Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.
- The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract.
- The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.
- Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

NOTICE AND DISCLAIMERS :

- © 2022 International Business Machines Corporation. No part of this document may be reproduced or transmitted in any form without written permission from IBM.
- U.S. Government Users Restricted Rights — use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.
- Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. This document is distributed “as is” without any warranty, either express or implied. In no event, shall IBM be liable for any damage arising from the use of this information, including but not limited to, loss of data, business interruption, loss of profit or loss of opportunity. IBM products and services are warranted per the terms and conditions of the agreements under which they are provided.
- IBM products are manufactured from new parts or new and used parts. In some cases, a product may not be new and may have been previously installed. Regardless, our warranty terms apply.”
- Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.

- Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.
- References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.
- Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.
- It is the customer’s responsibility to ensure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer’s business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer follows any law.

Agenda

- A look back at Db2 11.5.7
 - Summary
 - Deep Dives
- A peek at Db2 11.5.8
 - Summary
 - Deep Dives
- Reference:
 - Db2 11.5.0
 - Db2 11.5.4
 - Db2 11.5.5
 - Db2 11.5.6

Agenda

- **A look back at Db2 11.5.7**
 - Summary
 - Deep Dives
- A peek at Db2 11.5.8
 - Summary
 - Deep Dives
- Reference:
 - Db2 11.5.0
 - Db2 11.5.4
 - Db2 11.5.5
 - Db2 11.5.6

Db2 11.5.7 Highlights



Hybrid & Multi-cloud

Db2U operators automate deployment, upgrade and HA in Kubernetes environments.

Snapshot backup and restore. Seamlessly backup & restore all K8s resources in addition to customer data.

Db2U on Cp4D adds in Vault support to enhance security of credential and certification management



Application Development

Client support:

- Prefetch support for result set
- Increase fetch buffer size

Federation:

- Pushdown enhancements for:
 - CSE
 - Substring functions
 - JOIN and group-by
 - Varchar comparison
- FFNR support for Snowflake
- Fenced mode



Columnar Enhancements

- Improved Performance for Trickle (small) insert for all products.
- Real-time Statistics Collection (JITS)
- Schema Backup (Db2 W/IIAS)
 - TSM Support



Cognitive and Emerging

Db2 Graph supports PPCLE and Zlinux

New manage commands for graphs: `startSQLCapture`, `listSQLStatements`, `stopSQLCapture`, `validateDataIntegrity`

Enhancements to Graph Query Editor and Modelers

Db2 11.5.7 Highlights



Availability

- Log archiving to object store
- Zlib backup and log file compression
- HADR Hang detection for AIX
- Pacemaker/CoroSync
 - PPCLE support for 2-node
 - Software stack refresh
- DB24LUW-I-1048 RFE: fail restore if table space error occurs.



pureScale

- Z Linux Support (separate from 11.5.7 but also 4Q shipping on 11.5.1 level!)



Core Engine / SQL

- P10 support
- DB24LUW-I-1018: Avoid SPoF for Db2's Critical Files (aka Mirrored Log control files)
- Object Store Support Enhancements:
 - Upgrade to latest vendor SDK
 - Lift 5GB file size limit
- Rename table support when IOE exists on table – also adds AMT support
- Enable SNI in C-based clients

installFixpack now license acceptance 11.5.6 and on



License terms, third party and opensource products are changing in Db2 Modpack releases (as they contain new content).

Starting with Db2 11.5.6 installFixPack command will prompt users to accept the license terms.

The prompt for license agreement appears unless user passes a mandatory `-y` parameter in a non-interactive installation

Real Time Statistics Collections (RTS / JITS)

Prior to 11.5.7, if a table has no statistics or outdated statistics then:

- row organized: 5s of “synchronous” stats collection ,
 otherwise use “fabricated” (estimated stats based on meta data)
- column organized: only “fabricated” stats

As of 11.5.7, column organized now also supports synchronous stats collection – ON by default. You can: `db2set DB2_REDUCED_OPTIMIZATION="COL_RTS OFF"`)

If you see frequent timeouts in the stats log, consider:

- increasing frequency of manual collection
- change sampling rate for large tables via RUNSTATS profile
- increase time budget via opt profile or embedded guideline, e.g.:

```
<OPTGUIDELINES><RTS TIME='10000' /></OPTGUIDELINES>
```

-increase time budget via registry variable:

```
db2set DB2_ATM_CMD_LINE_ARGS="-rts-time 10000"
```

Knob to RESTORE when table spaces failures occur 1/2

DB24LUW-I-1048 - Add knob for restore to fail or succeed if an error with a table space is detected

Historically, Db2 recovery (restore, rollforward, HADR standby replay, crash recovery) tries to make available as much data as quickly as possible. As such, any error in one table space does not stop the recovery operation, it continues with as many tablespaces as it can.

Though this works well for getting a subset of data up quickly and w/o intervention, it has been requested to also allow the choice to fail recovery if any issue occurs -- “all or nothing approach”.

Knob to RESTORE when table spaces failures occur 2/2

In 11.5.4, we introduced reg var which impacted rollforward (including HADR):

`DB2_FAIL_RECOVERY_ON_TABLESPACE_ERROR`

NO: (Default) Previous behavior of not failing operation on table space error and table space state goes recovery pending

YES: Database/table space rollforward and HADR standby replay fails if any table space error is seen and table space state untouched

In 11.5.7:

Add restore to `DB2_FAIL_RECOVERY_ON_TABLESPACE_ERROR`.

If set to YES, if restore has an issue with a table space container then the restore will fail.

It will consider any error when restoring the table space definition/containers (including bad paths, can't allocate initial space, etc.)

Support for Archiving LOGs to REMOTE storage

DB24LUW-I-862 –Full support for log archiving to remote storage (AWS S3) ▪

Add new log archive method type DB2REMOTE; works just like DISK▪

Similar format to BACKUP and LOAD using storage access alias:

DB2REMOTE://<alias>//<storage-path>▪

Example:

```
UPDATE DB CFG FOR SAMPLE USING LOGARCHMETH1 DB2REMOTE://<alias>//<storage>
```

Cloud Object Store Enhancements

1/2

Two main enhancements:

- COS is updated to also support native vendor SDKs
- streaming and multi-part uploads (PUTs)
 - no longer require staging full file locally
 - avoid single file size limits

Overall `DB2_ENABLE_COS_SDK` can be used to enable / disable. Default = ON
Supported where COS is currently supported (Linux (x86 only))

When set to OFF, the legacy LibCurl approach is used (which is subject to known limitations, such as a 5GB individual file size, and lack of streaming multi-part upload support).

When set to ON, the new native SDK will be used. It does not support LOAD automatic decompression of .ZIP files nor LOAD CLIENT access to remote storage.

In the future (next new major release) plan is to support only the new SDKs.

Two ways to set “multi-part” size:

- 1) Db config parameter: MULTIPARTSIZEMB, **Default=100**, Range: 5 – 5120

- 2) Directly on the “commands” that support COS.
 - 1) LOAD ... COPY YES .. TO device/directory .. [MULTIPARTSIZEMB size]
 - 2) CREATE EXTERNAL TABLE command new “options” MULTIPARTSIZEMB=x

In all cases:

Specifies the size, in MB, of each part in a multi-part upload to object storage.

The specified part size directly impacts the maximum possible object size, since object storage vendors enforce a limit on the total number of parts in a multi-part upload.

DB24LUW-I-1018 Avoid Single point of failure for Db2's critical files

Ability to set up a secondary <path> to store duplicate copies of metadata/control files critical to Db2.

The files currently supported/maintained in the <path> are:

Table space
information files:

SQLSPCS.1,
SQLSPCS.2

Storage group
control files:

SQLSGF.1, SQLSGF.2

Buffer pool
information files:

SQLBP.1, SQLBP.2

Global log control
files:

SQLOGCTL.GLFH.1,
SQLOGCTL.GLFH.2

Log control files:

SQLOGCTL.LFH. 1,
SQLOGCTL.LFH.2

Mirrored logging
information file:
SQLOGMIR.LFH

To ENABLE, using a valid patch for <recovery-path>

```
db2 UPDATE DATABASE CFG FOR <database-alias>  
      USING ctrl_file_recov_path <recovery-path>
```

Reactivate the database

To DISABLE, set the path to NULL, DEFAULT or empty string

```
db2 UPDATE DATABASE CFG FOR <database-alias>  
      USING ctrl_file_recov_path <recovery-path>
```

Reactivate the database

Mirrored Control Files

3/3

Once enabled, every update to the “primary” files are “mirrored” to identical files in the <recovery-path>.

A failure to write to the <recovery-path> will have the same impact to the driving operation as if it was the <primary-path> IO that failed.
(Impact depends on operation + failure.)

If <recovery-path> is unavailable at first connect/activate, you need to DISABLE the feature – or fix the issue for the DB to start-up.

Agenda

- A look back at Db2 11.5.7
 - Summary
 - Deep Dives
- **A peek at Db2 11.5.8**
 - **Summary**
 - **Deep Dives**
- Reference:
 - Db2 11.5.0
 - Db2 11.5.4
 - Db2 11.5.5
 - Db2 11.5.6

Db2 11.5.8 Candidates



Hybrid & Multi-cloud

- Deploy on Amazon EKS
- Deploy on Azure AKS
- RHOS on AWS & Azure - ROSA and ARO
- Deploy on Rancher RKE
- Release on operator hub.io
- Integration with Velero for BAR
- Deploy on Google Cloud GKE



Application Development

- Multi-row Fetch - new regvar to force the count to match Z behavior
- Federation Enhancements:
 - Perf improvements in: FFNR, Pushdown for Impala JDBC, CSE
 - MongoDB – SSL, authentication and multi-host
 - SSL Support for HANA w/ ODBC
 - Add MySQL8
- Client Enhancements:
 - JCC improved connection pool
 - D2driver validation tool
 - Large result set query perf
 - .NET Visual Studio 2022 support
 - OpenJDL 17 runtime support



Columnar Enhancements

- LOAD enhancements
 - Performance Improvements
 - Applies page level varchar compression techniques (if enabled)
- Logical Columnar Backup (aka COISBAR/Schema Backup):
 - Make available as TechPreview on linuxAMD/ppcle across db2 installs (containers and on-prem)



Cognitive and Emerging

- GIU Enhancements to Query Screen
- 3 new tutorials

Db2 11.5.8 Candidates



Availability

- db2adult to upload load copy files to TSM
- LogReadAPI to map timestamp to LSN for (for logical rep)
- Log replay improvements (HADR/rollforward/crash recovery), and documentation for support recommended tuning parms
- GIU and silent installs now can install Pacemaker
- Improved deployment experience for Db2 HADR on AWS, Azure, GCP (automation and improved documentation)
- Pacemaker support for mutual failover on single node (aka single node active/passive w/ shared storage)



pureScale

- RHEL 8.6 support w/ RDMA on linuxamd+ppcle (via verbs)
- Simpler concurrent pureScale online modpack& fixpack update
- DIRECT_IO as default in Z-linux
- Periodic collection of RDMA roundtrip network stats
- pureScale on z15 support (may not be tied to modpack)
- pureScale TCPIP on AWS – deployment and support (may not be tied to modpack). Repeat for Azure.



Core Engine

- External table support to read from “binary numeric” format
- Additional logging for XML data (optional) when changed data capture is used.. Can be used by logical replication to replay XML w/o calling back to server.
- Force Load of GSKIT which comes with db2
- Extend db2RemStgManager to use storage aliases
- Monitor improvements: SSL info to be added to GET_MON_CONNECTION
- Catalog And Process Model Serviceability improvements
- db2support to automatically gather text search logs
- Stored procedures for formatted explain output
- Java, TSA, GPFS and Pacemaker updates
- Support for TLS 1.3
- GIU and silent installs now can install Pacemaker

Force loading of GSKit that came with Db2 driver

Currently, client-side GSKit loading attempts to load in the following order :

- 1) GSKit by unqualified name
- 2) GSKit from the fully qualified db2 path

The unqualified name can result in an unwanted GSKit level (usually most problematic if outdated).

Enable via ClientGSKitLocation=stock (Supported on all platforms (clients))

- db2dsdriver.cfg
 - db2cli writecfg add -parameter "ClientGSKitLocation=stock"
- db2cli.ini
 - db2cli writeini common ClientGSKitLocation stock

ClientGSKitLocation should **NEVER** be set to stock when the driver is embedded in some other IBM product which has already loaded GSKit.

Nothing is changed in the way the server loads GSKit.

TLS 1.3 support

Client/Server

- Server: Off by default in 11.5.8.0
- C Client – on by default
- Java Client – depends on JDK level, needs configuration
- HADR Primary-Standby
 - Receiver: Off by default in 11.5.8.0
 - New: Using `SSL_VERSIONS` dbm cfg along with client/server
 - Sender: on by default
- KMIP
 - On by default in 11.5.8.0
- New restrictions on certificates when using TLS 1.3
 - No SHA1 or SHA224 signatures
 - No RSA keysize < 2048

Old Clients that do not support 1.2 will not be able to connect when 1.3 is enabled (even if `SSL_VERSIONS` has `TLSV1` included)

- **Prior to 9.7 fp9, 10.1 fp 4, 10.5 fp3**

MON_GET_CONNECTION

Adding the following metrics:

- TLS Version
- Cipher spec negotiated
- Label name and fingerprint when connection was established

db2pd “-applications “ will also include this additional info

pureScale Concurrent Online Update

Automation of previous stages allowing for concurrent (in parallel) updates.

3 new options to installfixpack:

-autoupdate

updates entire cluster, Db2 automate which hosts to update concurrently while maintaining quorum.

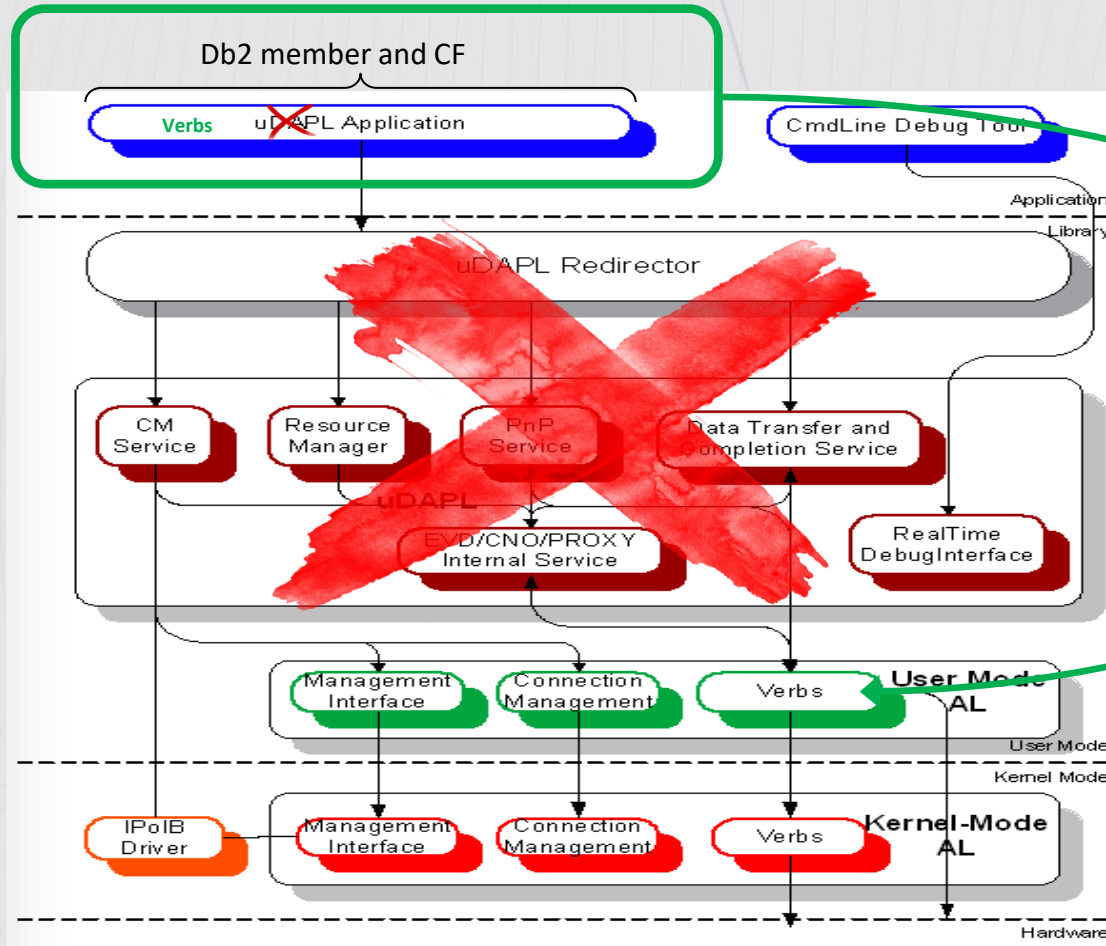

-concurrency <#_of_hosts>

specifies a maximum concurrency level (limit loss of processing power)

-show_update_order

use with above options; simply prints the sets and order it *would* run.

RDMA support on RHEL 8.x



Motivation
RHEL 8.1 removes support of uDAPL

Modification to Db2 and CF to replace uDAPL usage with Verbs directly

- Benefits:**
- Supported on RHEL 8.x
 - Reduce 3rd party code dependency
 - Potential performance improvement in reducing lock contention which can improve overall throughput

Load copy file support for TSM

1/2

DB24LUW-I-1567 Upload db2 load copy images with db2adutl from filesystem to tsm

Objective: ensure that all images that may be needed for recovery (both restore and rollforward) can be in TSM, and that they are used for recovery.

Part 1: db2adutl updates to support copying files to TSM

New syntax for db2adutl. (note: IMAGE will be undocumented but still works)

```
UPLOAD [[ AND REMOVE ]]  
  [ LOADCOPY [[ TAKEN AT <timestamp> ]][ <filename> ] ] ] ]  
  [ BACKUP [[ TAKEN AT <timestamp> ]][ WITH LOGS ] ] ] [ <filename> ] ] ]  
  [ LOGS [[ BETWEEN <Sn1> AND <Sn2> ]][ OLDER THAN <Sn1> ] ] [ CHAIN <n> ] [ FROM [ LOGARCHMETH1 | LOGARCHMETH2 ] ] ]  
  [ MGMTCLASS <mgmtclass> ]  
  [ DB2USER <db2 username> ] [ DB2PASSWORD <db2 password> ] ]
```

Files will be copied (removed if requested), history file updated (if possible).

Load copy file support for TSM

2/2

Part 2: Use of TSM images by recovery

Db2 currently uses the location written to the corresponding log record by default. The db2adult upload procedure is only capable of physically uploading the image and updating the history entry.

To ensure ROLLFORWARD can find the copy images, load copy recovery will be enhanced to look in multiple locations in order to find the load copy image. If set, Db2 override registry variable DB2LOADREC. The hierarchy will be:

- Location in log record

- Location in history entry

- TSM if configured

Reminder from 11.5.4

Columnar Table New Compression Algorithms

2 **NEW** page-level compression methods for string data:

- 'repeating pattern'-based compression (LZ4 based)

 - Work's well for: Geospatial data, URLs, Comment fields, etc.

 - Compression rate typically 2-4x but depends on frequency & length of patterns

- nibble based compression

 - 2x compression for strings with ≤ 16 unique characters

 - Works well for : phone numbers, coordinates, \$ amounts, etc

Stored compressed on disk and in buffer pool

Both of these **are off by default** (for on-prem) and require a registry variable to enable. Once set, any tables using these new compression mechanisms will not be readable by previous mod packs of Db2 11.5

Back to 11.5.8

Columnar Table - LOAD Improvements

1/3

LOAD will reuse code from the bulk insert that introduced in 11.5 GA. This itself will simplify and streamline LOAD. This will be on by default.

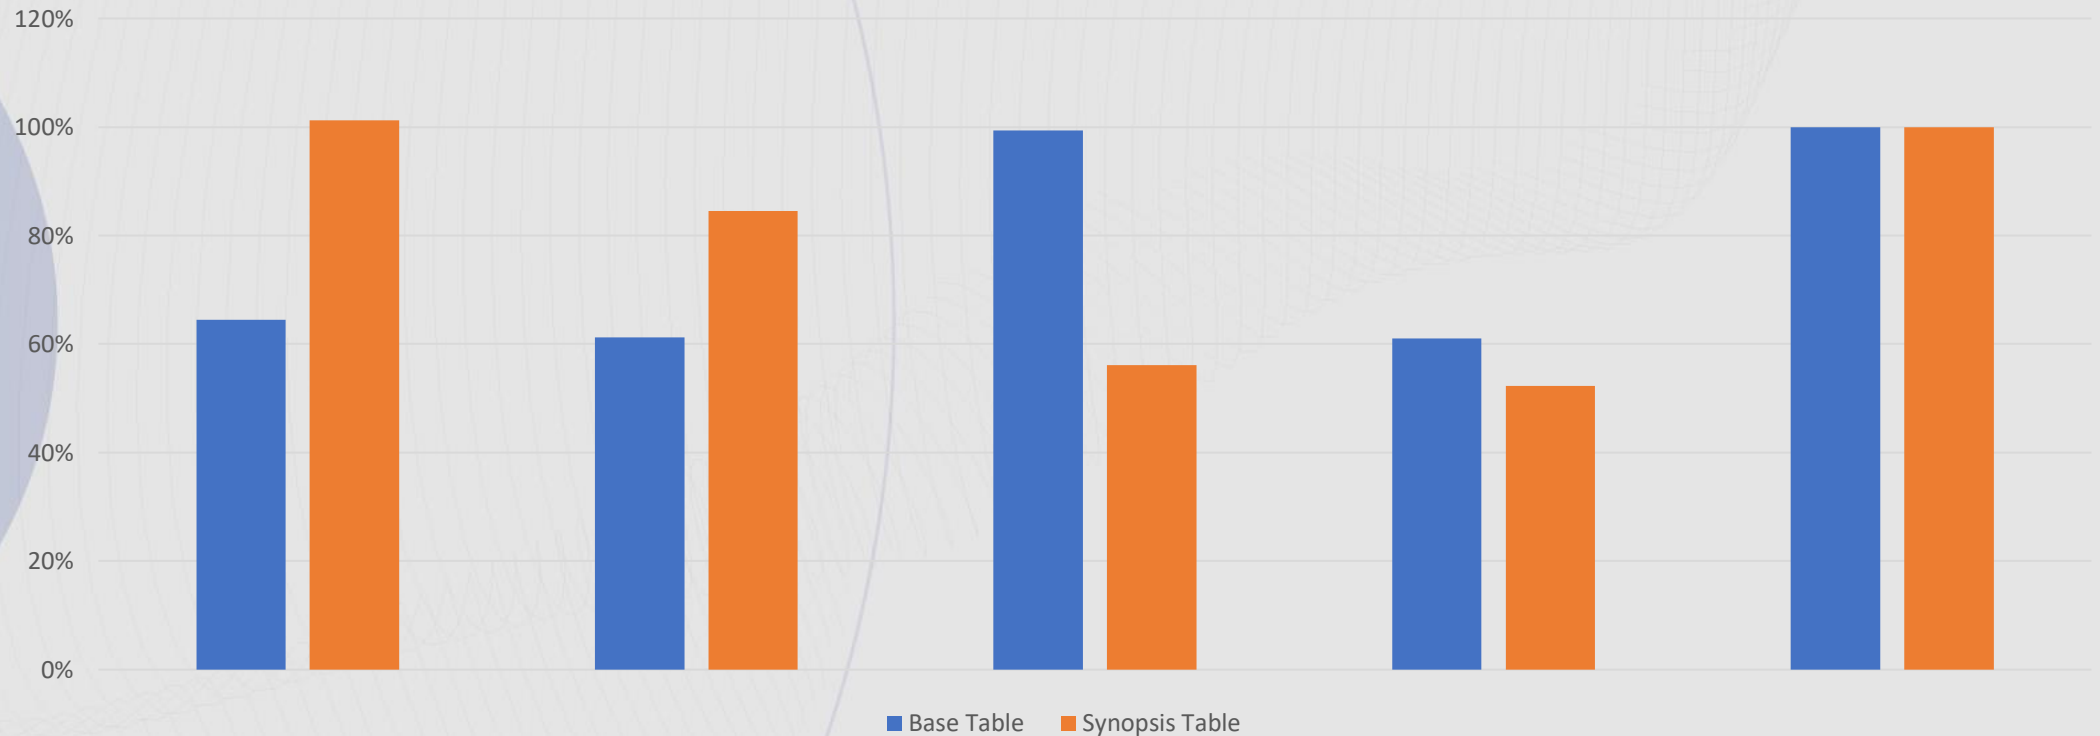
Secondly, the reuse allows the 2 new compression algorithms to be used (if enabled). This may result in better compression which itself can improve performance even more (less IO).

In lab tests, show perf improvements of 0-15% -- depends on data types in tables, how much it compresses, disk speed/load, etc.

Columnar LOAD enhancements – Size reduction

2/3

Table Size for 5 Different Tables Compared to Disabled (100% = same size)




* In lab results on early code drop. With enabled, compared to w/o enabled (as 100%)

Columnar LOAD enhancements – Speed

3/3

Elapsed LOAD Time for 5 Tables



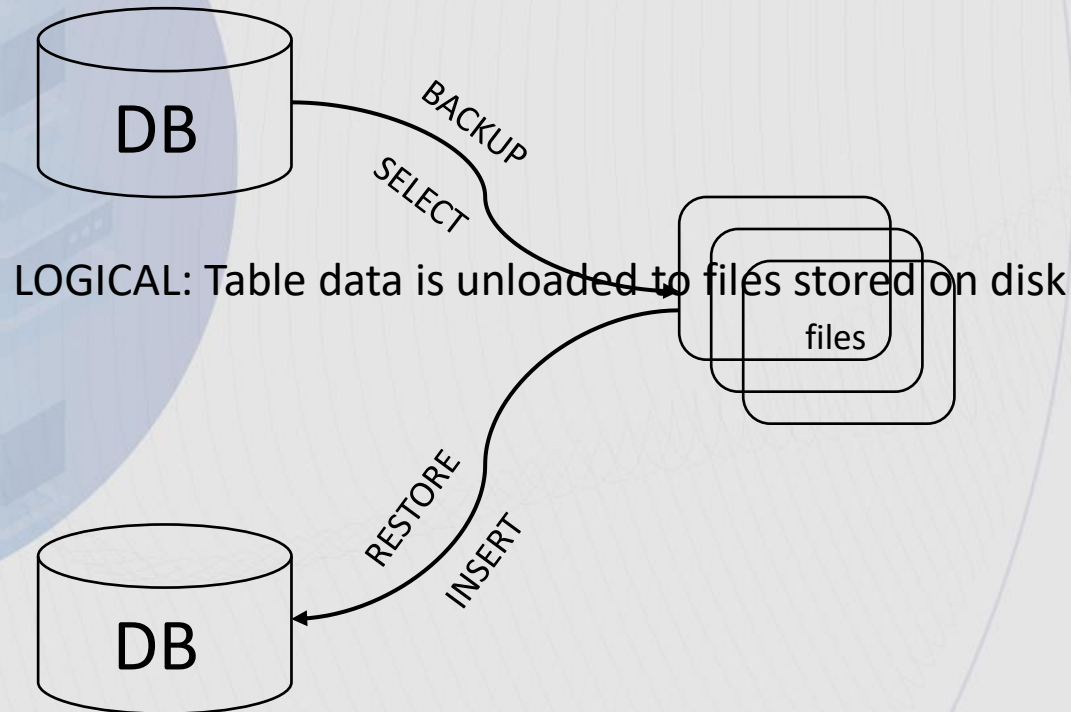
* In lab results on early code drop, w/ and w/o new feature

Logical Backup and Restore

(Previously known as Schema BnR, COISBAR.)

(Tech Preview outside IIAS)

What is Logical Schema BnR?



- **Columnar Tables only**
- Agnostic to instance setup (DPF/EE)
- Cross-platform (linuxamd & ppcle to start)
- Schema granularity on backup
- Table level granularity on restore
- Subject to all SQL rules: WLM, config
- R/W access to tables during backup
- Can also be used for easy table/schema copy from production to test

Agenda

- A look back at Db2 11.5.7
 - Summary
 - Deep Dives
- A peek at Db2 11.5.8
 - Summary
 - Deep Dives
- **Reference:**
 - **Db2 11.5.0**
 - **Db2 11.5.4**
 - **Db2 11.5.5**
 - **Db2 11.5.6**

Db2 11.5 GA Highlights



Cognitive and Emerging

- ML Optimizer Tech Preview
- Federation support for Block Chain Tech Preview



Hybrid & Multi-cloud

Docker container



Application Development

- Net core 2.2 support
- GO language Driver
- IDE for Visual Studio Code – a db2 extension
- Augmented Data Explored (ADE) beta



Columnar Enhancements

- Next Gen Bulk Insert for columnar tables
 - Vectorized Insert (and Update)
 - Reduced UNDO logging
- Update+Delete performance enhancements
- Automatic Dictionary Creation (ADC)
- Automatic REORG RECOMPRESS (compression of uncompressed data used to build dictionary)
- Vectorized ADC
- Support for LOB data type in columnar tables
- Numerous improvements to columnar query performance.

Db2 11.5 GA Highlights



Availability

- Advanced Log Space Management Tech Preview
- Parallel Logging when using Mirrored Logs



- Improved Cluster-wide free space management
- host-based firewall support
- Automate setup for public Ethernet monitoring
- increase ports, XI connection and worker limits
- OLIC and extent reclaim on by default
- Currently Committed across pureScale members
- 2x faster LOAD w/ Range Partitioned Tables



Core Engine

- External Table Support
- 4k Sector Support
- Auto Column Group Stats (CGS)
- Numerous additional monitor elements
- DROP/CREATE TABLE [IF EXISTS]
- Create Table As (CTAS) enhancements
- DBMS_APPLICATION_INFO and UTL_RAW package support (Oracle Compat)
- PLSQL Enhancements
- WLM:
 - Simplified thresholds
 - Cascaded drop service class

Db2 11.5.4 Highlights



Graph database

Graph database + Relational database in one store.

Tech preview for 11.5.4

Overlay a graph on top of Db2 data, and query Db2 via Gremlin language.



Hybrid & Multi-cloud

Db2 container for Red Hat OpenShift



Application Development

- Drivers and NoSQL wire listener updates
- Db2 REST APIs
- Python UDFs
- Programming language updates:
 - .Net 2.2,
 - .Net drivers for MS Azure,
 - Node JS v12,
 - Ruby/Rails latest,
 - Django 2.0



Columnar & MPP Enhancements

- Automatic INDOUBT Resolution for DPF
- Adaptive workload management for WORKLOAD=ANALYTICS (on prem now supported)
 - Plus session priorities,CPU Controls, sort threshold
- Columnar Page based VARCHAR compression
- Delayed synopsis table population
- Query Perf (early aggregation, early distinct, full outer join, join residual predicate support, NULL=NULL)
- Memory Stability (Compact VARCHAR)
- Truncate Table -- rollback support
- RID Scaler Function for Columnar/DPF
- Alter Table add column support for LOB

Db2 11.5.4 Highlights



Availability

- Automated HADR with Pacemaker for RHEL- Tech Preview
- Advanced Log Space Management
- Ability to block reorg pending operations through registry variable
- New column in MON_GET_HADR
HADR_LAST_TAKEOVER_TIME



Core Engine

- Optimizer Version Control
- Faster Database Activation
- Skipped locked Data For Queries
- Nested WITH support
- ALIAS support in WHERE clause
- Faster Index Splitting at non-leaf levels under high contention
- UTL_RAW, DBMS_LOCK, DBMS_STAT package support (Oracle Compat)



Security

- SSO with JWT Token
- Authentication Caching
- Security-Enhanced (SE) Linux support on RHEL 7 & 8
- Allow SSL_SVR_LABEL to be changed online



Cognitive and Emerging

- ML Optimizer - Tech Preview 2
- In Database Analytic functions
- Spatial Analytics
- Federation:
 - Parallelism
 - Numerous JDBC connectors

Db2 11.5.5 Highlights



Cognitive and Emerging

- ML Optimizer Tech Preview:
 - Support for additional queries



Hybrid & Multi-cloud

RHOS Container Updates



Application Development

- Updated Visual Studio Code IDE
- Federation Enhancements:
 - Even faster parallelism
 - MySQL CE ODBC
 - Hive 3.0 ODBC & JDBC
 - Snowflake ODBC
 - PostgreSQL V10 ODBC & JDBC
- Client Enhancements:
 - Odata – MQT and View support
 - JCC runtime support for OpenJDK13



Columnar Enhancements

- Reorg Recompress
 - Performance and Concurrency Improvements
 - Applies page level varchar compression techniques
- Schema Backup (Db2 W/IIAS):
 - Reorg Table reclaim support for enabled tables
 - RCAC Support
- Query Performance and Memory reduction. For Group By and Hash Joins with VARCHARs

Db2 11.5.5 Highlights



Availability

- GA: Pacemaker/CoroSync Support for Automated HA (HADR)
- Advanced Log Space Management:
 - Support for Mirrored Logs



pureScale

- Faster pureScale online modpack& fixpack update (concurrent)
- Automatic Cluster Validation (periodic)
- Lightspeed RDMA ping



Core Engine

- Spatial Analytics:
 - Support for Shape Info in db2se
 - Performance improvements
- Schema Level Authorization
 - Supported on all platforms
 - Audit support added
- db2histmon – add mon_report style functions
- JWT
 - Support for multiple labels in config
 - db2pd Enhancements

Db2 11.5.6 Highlights



Cognitive and Emerging

- Db2 Graph GA – including new UI
- ML Optimizer tech preview



Application Development

- R support for UDX
- CLI/ODBC Result set prefetching
- Federation Enhancements:
 - DRDA bulk insert improvements
 - Postgres 12 support
 - Faster Spark JDBC support



Columnar Enhancements

- Improved Performance for Trickle (small) insert – Db2 Wh products only
- Columnar Temporal Tables
- Logical BNR (Db2 Wh products only)
 - Progress monitor
 - LOAD with exception tables
 - RCAC support
 - WLM integration
- Query Performance and Memory
 - CVC - Aggregation DISTINCT
 - Better selection/usage of OTF
 - Join – overshoot reduction
- Improved Performance with Updates
 - Improved Pushdown
 - Select after update



Hybrid & Multi-cloud

RHOS Container Updates

- RHOS 4.6 support
- Multiple standby support for HADR
- Updated version levels for storage



Install / Deployment

- Integrated install for Pacemaker and corosync
- Mandatory prompt for license agreement
- db2updserv discontinuation

Db2 11.5.6 Highlights

Availability

- Pacemaker/CoroSync
 - Integrated bundling and Install
 - Alternate Quorum and VIP config
 - Improved PD (including updates to Pacemaker OS)
 - Enhanced Resiliency
- NX842 hardware compression with encryption for backup/restore
- HADR – support multiple standbys in K*'s environments (NAT)
- HADR Hang Detection (Linux only)
- Advanced Log Space Management, additional support for:
 - Online Backup and rollforward
 - HADR
- db2adutl delete performance enhancements

pureScale

- Diag.log contention “protection”,

Core Engine / SQL

- >1012 column support
- Db2fmtlog externalize basic formatting
- External Table - support for hidden columns
- Adaptive WLM for all Configurations
 - Integrated CPU Controls
 - Sort query memory scale-back
- Increase Parallelism/Performance For Set Integrity
- CLP Hashing Performance Improvements
- ADMIN_MOVE_TABLE
 - Perf improvements for BLU via parallel insert (4x)
 - Support for generated always columns
- Spatial Analytics:
 - KML Import,
 - GeoJSON Inport/Export
- Updates to default port numbers to avoid Linux Ephemeral port
- LOCK NO WAIT and WAIT time on Select/update/delete