Db2 V11.5 Overview & Latest From the Lab

Keri Romanufa
IBM

Email: keri@ca.ibm.com
LinkedIn: keri-romanufa
Twitter: @RomanufaKeri

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 :

Agenda

- A look back to Db2 11.5.0
- A look back at Db2 11.5.4
- A lock back at Db1 11.5.5
- A peek at some of the future 11.5.x+ candidates

Agenda

- **A look back to Db2 11.5.0**
- A look back at Db2 11.5.4
- A look back at Db2 11.5.5
- A peek at some of the future 11.5.x+ candidates

Db2 11.5: What's Up With the Numbering?

Db2 11.5 follows a continuous delivery model   (same as 11.1)

If a delivery vehicle has:

- new function -- will increment Mod level   (same as 11.1)
  - But it will also reset the FP level to 0   (*NEW*)
    → 11.5.M.0

- ONLY fixes --   will increment just the FP level   (same as 11.1)
  - This addresses the issue where 11.1 "ifixes" were not easily distinguishable from the base mod pack nor other ifixes off the same mod pack.
    → 11.5.M.1 and higher

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

⚙️⏱️ **pureScale**

- pureScale Improved Cluster-wide free space management

- pureScale host-based firewall support

- Automate setup for public Ethernet monitoring

- pureScale increase ports, XI connection and worker limits

- pureScale OLIC and extent reclaim on by default

- Currently Committed across pureScale members

- PureScale 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

Agenda

- A look back to Db2 11.5.0
  - **A deep dive into some 11.5.0 items**
- A look back at Db2 11.5.4
- A look back at Db2 11.5.5
- A peek at some of the future 11.5.x+ candidates

Support for storage devices using 4K sector size

This support can be enabled by setting the DB2_4K_DEVICE_SUPPORT registry variable to ON

Allows Db2 databases to be created on devices that only allow 4k IOs.

Restrictions and limitations:
- The use of DMS Raw containers is not supported
- Backup and load copy files will be slightly larger
- There may be a small performance penalty accessing:
  - Small non-inlined LOBs
  - Backup or load copy files created prior to the enablement of 4K device support

Columnar Insert/Update/Delete Performance

- Db2 11.5 greatly expands core-friendly parallelism for SQL-based IUD operations on columnar tables
    - KIWI: Kill It With Iron   (Slide stolen from Chris D! Only fruit acronym I have!)
    - Maximize CPU cache, cache-line efficiency

- Critical to maximize ETL/ELT batch performance

- Many general improvements, but primary focus on bulk operations

Performance - ETL Performance Example

- Data ingest rate
  - 1 TB/hour before enhancements
  - *Now ~5 TB/hour (IIAS)*

- >10 TB data

- Table remains online

- Combined features (in Db2 W today)
  - GA - ET load
  - future - Columnar Reduced Logging
  - GA - Columnar Parallel insert
  - GA - Columnar Vectorized insert
  - GA - Columnar Optimized bulk insert codepath

Parallel Insert Degree/Time/Speedup
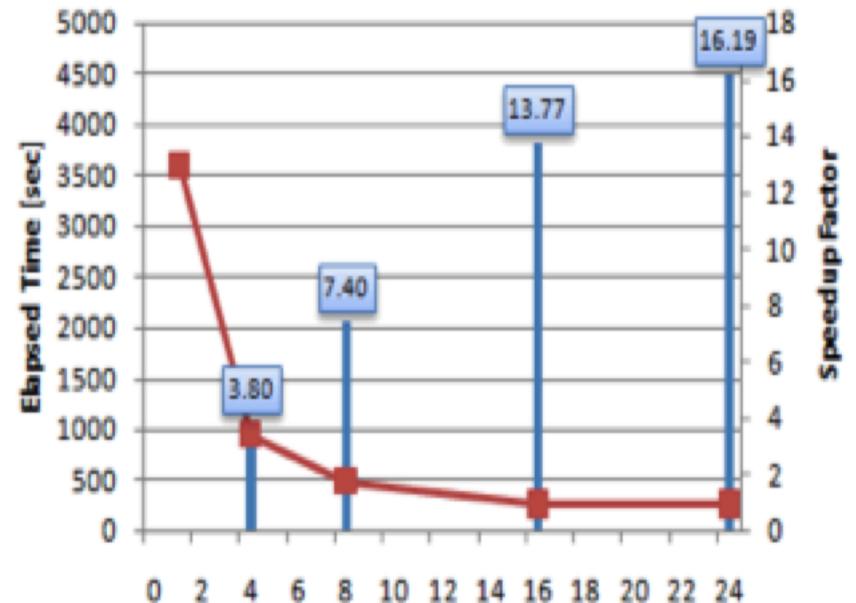
External Tables

A simple mechanism that lets you access external "files" within an SQL statement just like a table
- Can also be used to load from or unload to external files
- Can be used to define a permanent external table or directly within a SQL statement

Currently supports CSV/delimited files and local/remote sources

Example:
```
create external table ext_orders(order_num INT, order_dt TIMESTAMP)
                              USING(dataobject('/tmp/order.tbl') DELIMITER '|');

insert into orders (select * from ext_orders);
```

Example:
```
insert into orders (select * from external '/tmp/orders.txt' using(REMOUTESOURCE
GZIP delimiter ','));
```

Db2 pureScale Cross-member Currently Committed 1/2

In pureScale environments, the Currently Committed isolation will avoid lock wait when:

- Concurrent access allows a member to use a older (but still valid) version of a page that does not yet show a IUD operation

- The application performing the row read, and the application performing the row UPDATE or DELETE, reside on the same member

- The application performing the row read, and the application performing row INSERTs, reside on any member

Db2 pureScale Cross-member Currently Committed   2/2

In current pureScale environments, the Currently Committed isolation method can still result in a wait if:

- A row-reader finds the row is being updated or deleted (locked) by an application on a different member

This feature brings completeness to Currently Committed behaviour in pureScale environments, by allowing a row-reader to go out to the another member to return the currently committed data from the other members logs.

GoLang for Db2

- Now available @ https://github.com/ibmdb/go_ibm_db

How to run sample program

example1.go:-

```
package main

import (
    _ "github.com/ibmdb/go_ibm_db"
    "database/sql"
    "fmt"
)

func main(){
    con:="HOSTNAME=host;DATABASE=name;PORT=number;UID=username;PWD=password"
    db, err:=sql.Open("go_ibm_db", con)
```

- https://github.com/ibmdb/ - contains 8 open source drivers (Go, Python, Jupyter Notebook, Sequelize, PHP, Java (Spring), Node.js, Ruby)

16

Agenda

- A look back to Db2 11.5.0
- **A look back at Db2 11.5.4**
- A look back at Db2 11.5.5
- A peek at some of the future 11.5.x+ candidates

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

Roadmap: https://ibm-analytics-roadmaps.mybluemix.net/

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

Agenda

- A look back to Db2 11.5.0
- A look back at Db2 11.5.4
  - **A deep dive into some 11.5.4 items**
- A look back at Db2 11.5.5
- A peek at some of the future 11.5.x+ candidates

Security -- Authentication cache   1/3

Intended to relieve performance impact due to bottleneck on authentication backend
- Workloads with extremely short duration connections which occur repeatedly using the same, limited set of authorization IDs

Applies to all password authentication plugins supported by Db2

Cache is configured based on maximum number of users to be cached and duration of time for cache (default is 3 minutes)
- Credentials are not cached

Security -- Authentication cache   2/3

Without a cache – every connection needs to call out to LDAP Server
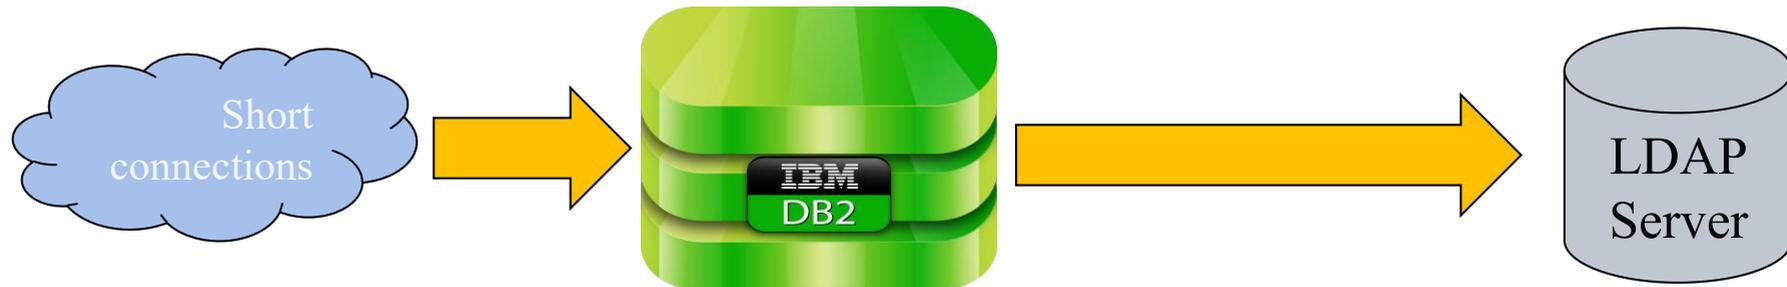
Security -- Authentication cache   2/3

With a cache – only a small subset need to, rest are serviced by the cache

Security -- Authentication cache   3/3

Controlled by 2 configuration parameters:

AUTHN_CACHE_USERS –
- Number of entries to be kept in the Db2 Authentication cache.
- Default 0   (==OFF)

AUTHN_CACHE_DURATION –
- Time in minutes for which an entry is considered valid and available for reuse in the Db2 Authentication cache.
- Default 3min

New SQL statement, to allow SECADM and DBADM to clear the cache:
- FLUSH AUTHENTICATION CACHE

GET_MON_DB also updated to be able to monitor the efficiency of the cache

OPTIMIZER VERSION CONTROL 1/2

Goal:

Allow a DBA to quickly revert the optimizer behavior in terms of:

- Query rewrites/transformations
- Access plan generation

to match a previous version/release.

Single control* that can quickly be used in emergency situations where a regression(s) may be experienced after an upgrade or update.   Can also be used pro-actively.

It is not recommended to be run with long-term.

*Prior to this, a large number of both documented and undocumented reg vars would need to be set to revert behavior.

OPTIMIZER VERSION CONTROL   2/2

New registry variable:

DB2_OPTIMIZER_VERSION

Default = current version

Can be set to any 4 part release name since 10.5.0.0

E.g 11.5.0.0 to match 11.5 GA

Does not prevent new plans due to change in statistics or other non-optimizer code changes.

This is applied first,   followed by any other SQL compiler registry variables which can then override it's underlying settings.

Availability -- FASTER Database STARTUP

ON by DEFAULT

Speeds database startup by
 up to 120x.

The larger the buffer pools and
lock list the larger the speed-up.

Database Activation Time

43x faster

120x faster

11.5   11.5 Tech Preview   11.5 future

■ LINUX - 320 GB   ■ AIX 512 GB
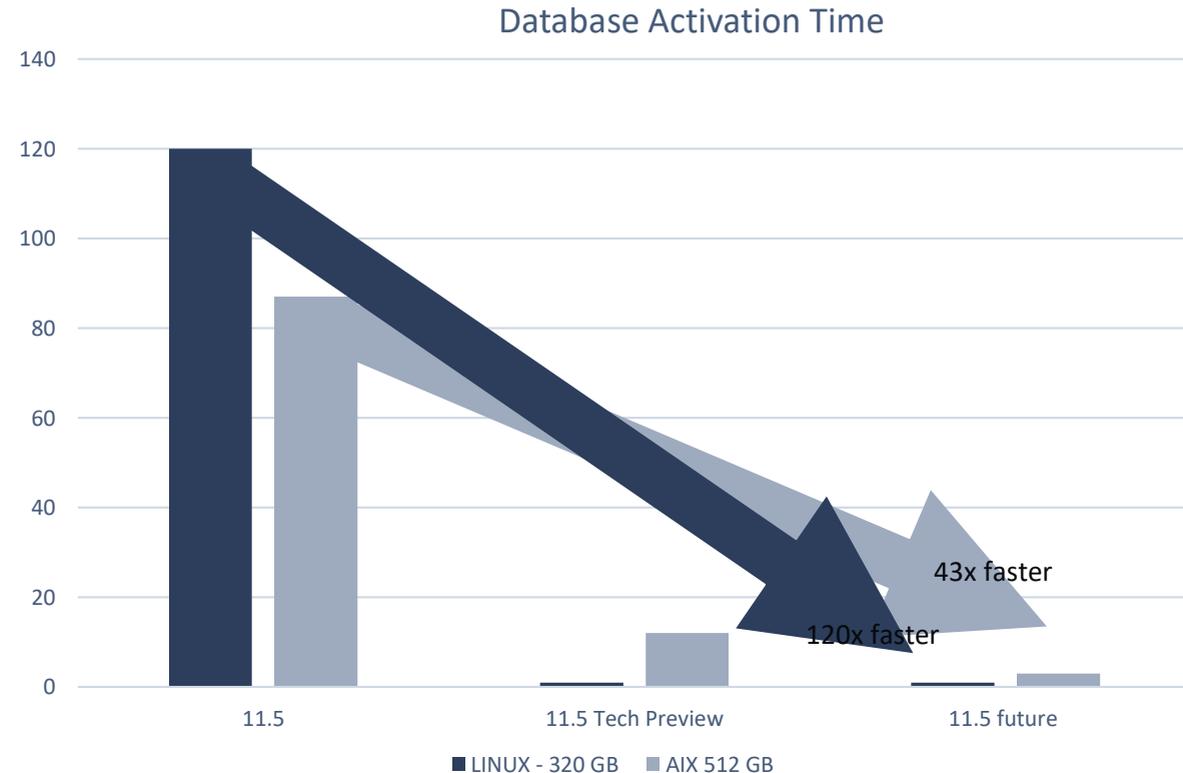
Storage -- Columnar Compression IMPROVEMENTS   1/2

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** 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

Storage -- Columnar Compression IMPROVEMENTS   2/2

Internal Lab Results based on customer-like data.

Customer-like tables with VARCHAR and other column types

Table Sizes – Financial Data
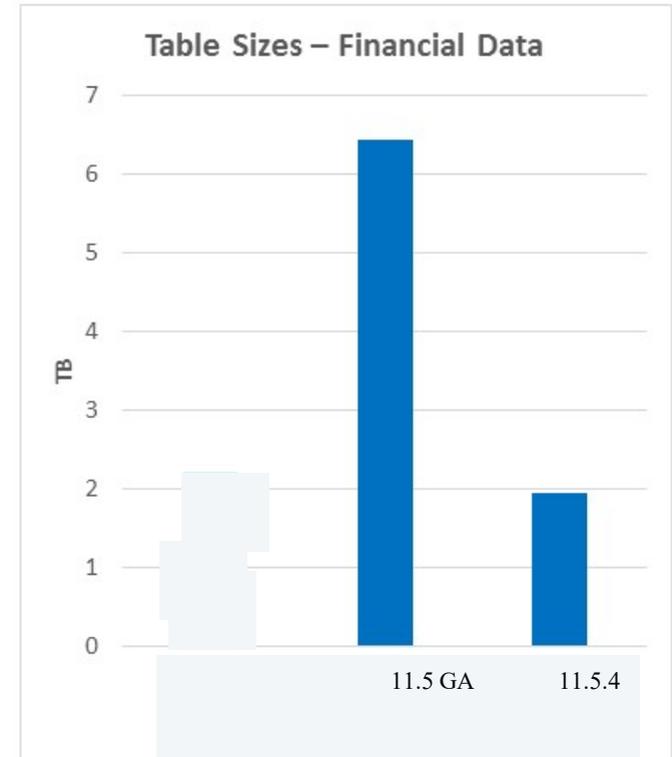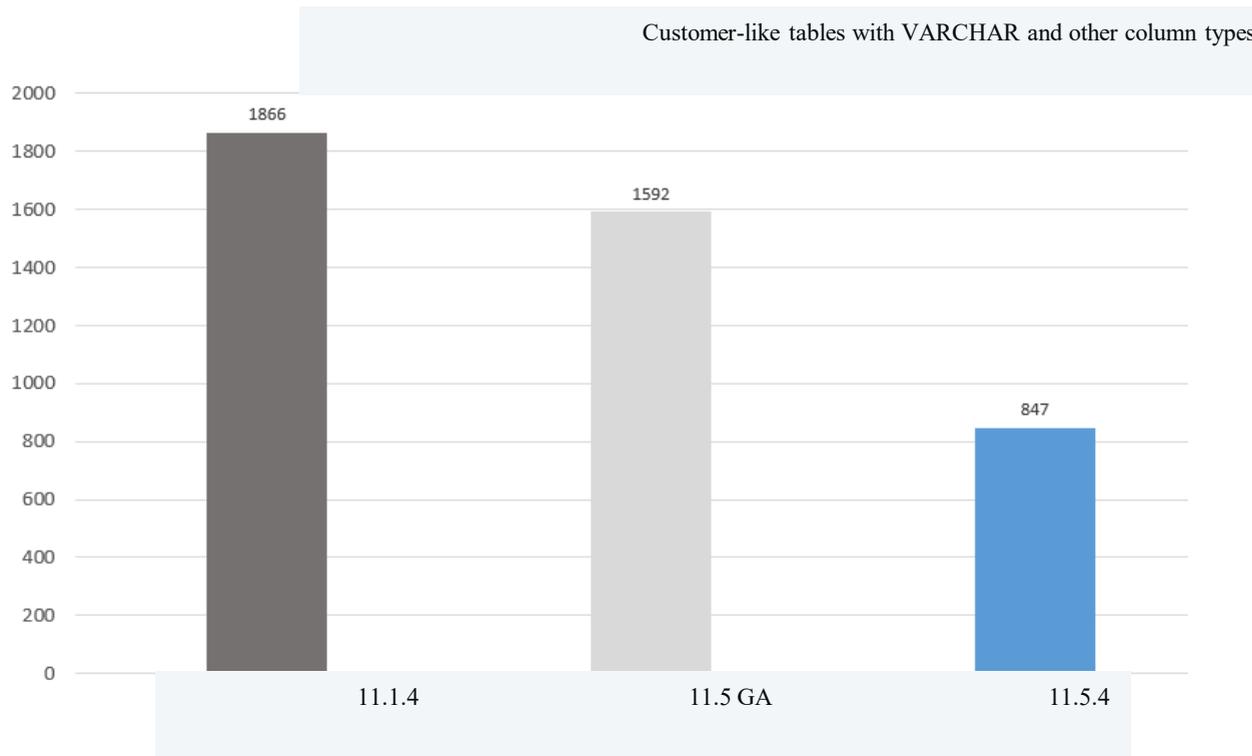
Simplicity-- Adaptive WLM

- AWLM will have a staged delivery to Db2
  - 11.5.4: AWLM available for use by analytics (Columnar) customers
    - Available to customers who set DB2_WORKLOAD = ANALYTICS

  - Stage 2: AWLM available for use by all Db2 customers

- Moving to AWLM will require an outage to allow the implementation of a new default system configuration
  - A new "opt-in" procedure will be provided

Simplicity-- Advanced Log Space Management   1/2

As part of a focus on simplifying log management, the first delivery will result from a focus on avoiding "log full" scenarios caused by long running, low volume transactions hold up log space

E.g. transactions that span multiple log files with little content in the intermediate files

Log records from long running transaction

Spanned logs

S0000000.LOG   S0000001.LOG   S0000002.LOG   S0000003.LOG

Active log files

31
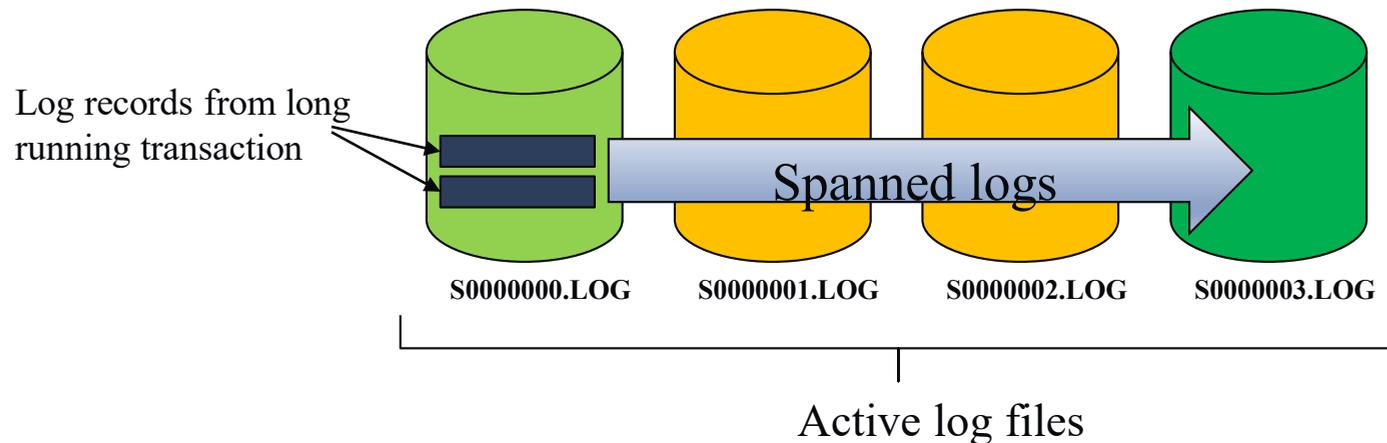
Simplicity-- Advanced Log Space Management   2/2

Solution is to extract log records for long running active
transactions to a separate file and allow intermediate log files to
be closed, archived, and reused

Log records from long
running transaction

TXID_EXTRACT.LOG

S0000000.LOG   S0000001.LOG   S0000002.LOG   S0000003.LOG

Available log files   Active log files
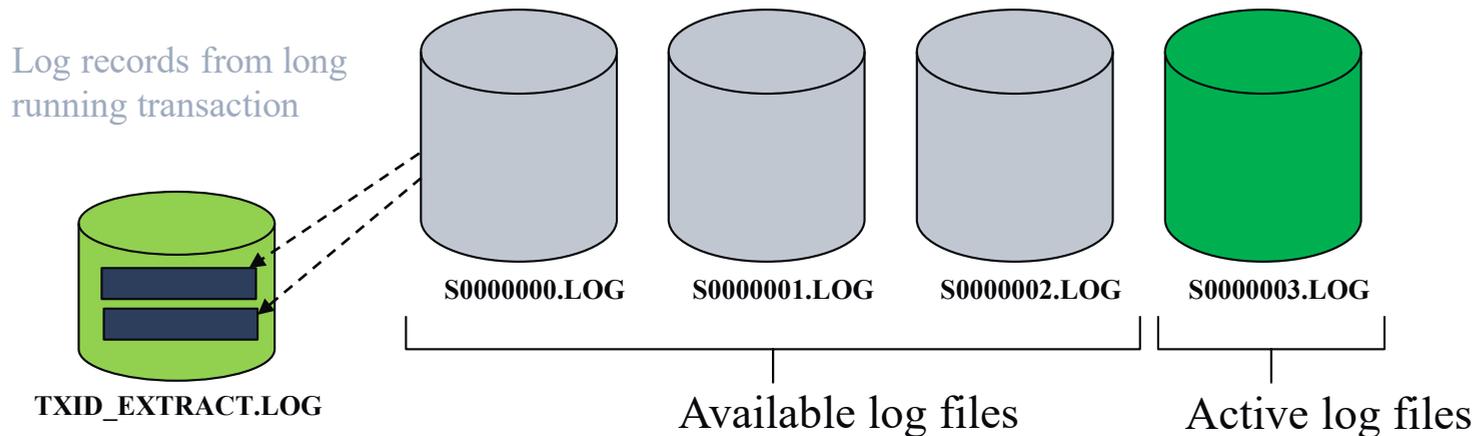
Db2 Graph - Tech Preview

Your data remains in Db2

- Continue to read and make updates to your data, without any impact on performance

- Leave your applications as they are

Simply "overlay" a virtual property graph over your Db2 data

- Uses the referential information to automatically create the Graph schema

- You can map tables or create views to map them into graph vertices and edges

- Add more relationships by editing the Graph schema, as long as the underlying data exists in Db2
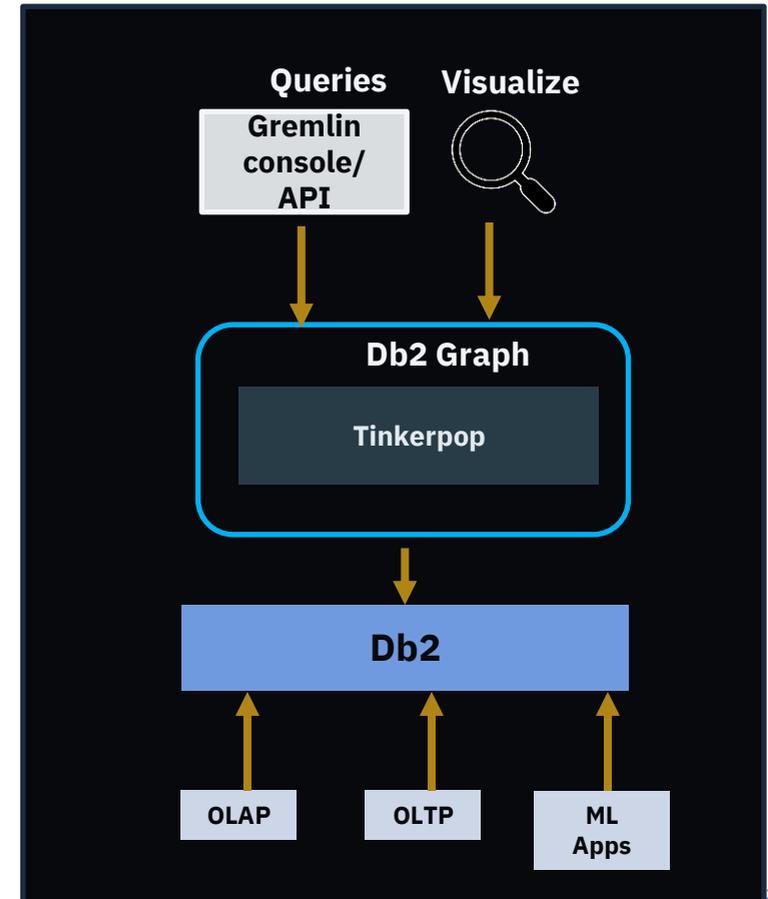
33

Agenda

- A look back to Db2 11.5.0
- A look back at Db2 11.5.4
- A look back at Db2 11.5.5
- A peek at some of the future 11.5.x+ candidates

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 OBDC
  - PostgresSQl 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

Agenda

- A look back to Db2 11.5.0

- A look back at Db2 11.5.4

- A look back at Db2 11.5.5
  - **A deep dive into some 11.5.5 items**

- A peek at some of the future 11.5.x+ candidates

Compact Varchar Phase 2

Impact

- Improved memory efficiency for wide VARCHARs in CDE Group By and Join queries
- Reduce memory consumption, spill I/O and OOMs primary impact focus
- Performance improvements
- Increase in concurrency within Group By and Join operator

Results

- Overall workload elapse time, memory footprint and spilling greatly improved
  - Performance: Up to 2.9X overall workload, 17.6X individual query
  - Memory reduction: Up to 1.1X overall workload, 2.5X individual query
  - Spilling reduction: Up to 5.6X overall workload, >1200X individual query

Schema Level Security

Database

Administration : DBADM
Security : SECADM, ACCESSCTRL
Access : DATAACESS
Load : LOAD

Schema

Administration : SCHEMAADM
Security : ACCESSCTRL
Access : DATAACESS, *SELECTIN, INSERTIN, UPDATEIN, DELETEIN*
Load : LOAD

Object

Administration : Owner/CONTROL
Security : Owner/CONTROL
Access : *SELECT, INSERT, UPDATE, DELETE*
Load : LOAD
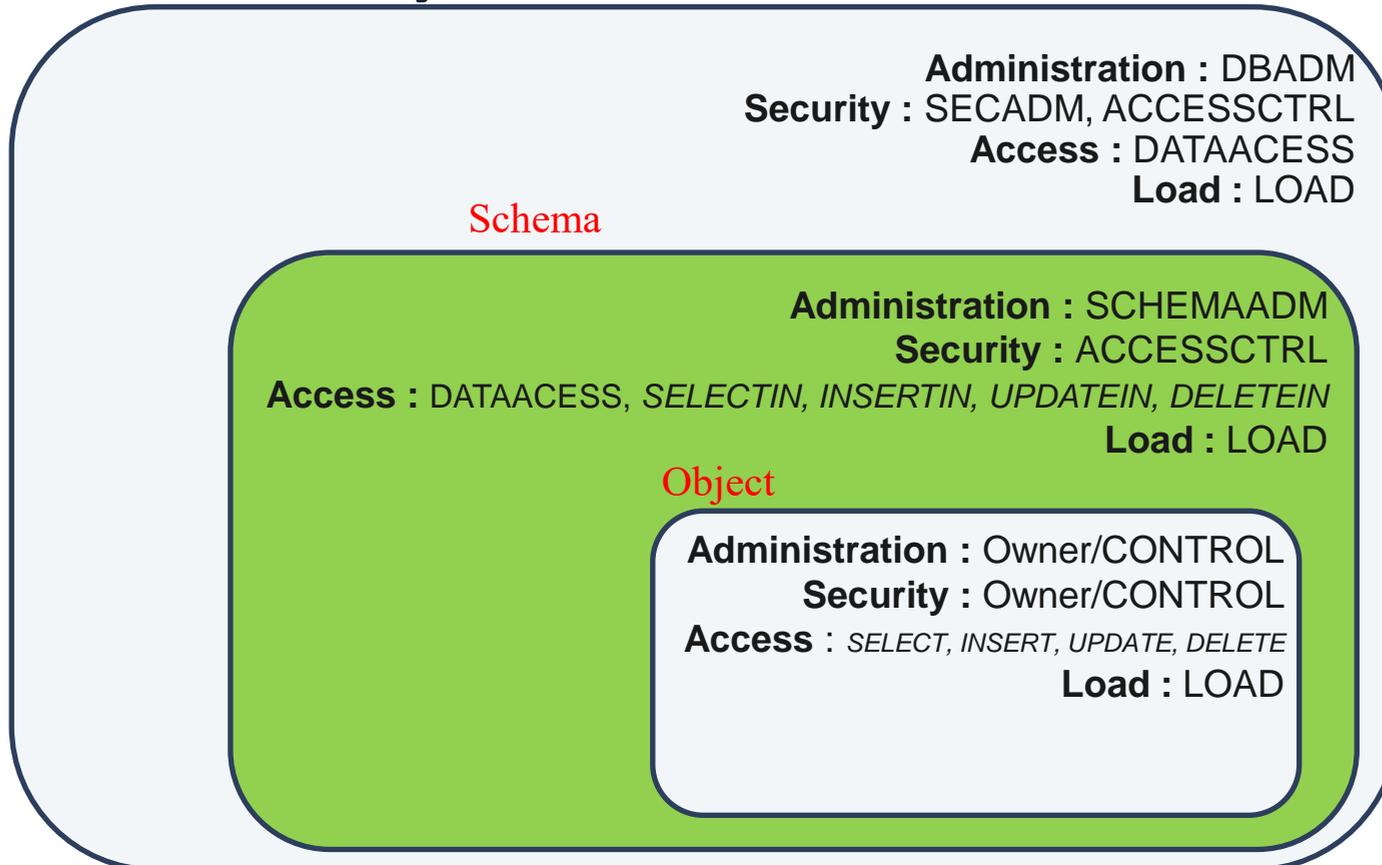
AUTHORITY

PRIVILEGE

SECURITY

Pacemaker & Corosync   1/3

2 Node automated HADR failover for Linux

V11.5.4.0 Technical Preview

- Cluster manager-aware integrated Db2 commands

- Integrated data collection via db2support
- New cluster manager configuration utility – db2cm
- Multiple instances & databases support
- Enhanced quorum type support with QDevice
- RHEL 8.1, SLES 15 SP1 support on Intel and Linux on IBM Z
- Validated on AWS with RHEL 8.1

V11.5.5.0 Technical Preview → GA

- Multiple Standby Support
- Fast redeployment via import & export support
- Two node support with fencing on AWS

Future GAs

- Integrated bundling and install
- Support on POWER (AIX and PPCLE)
- Active-Passive HA Configuration
- Mount point monitoring
- DPF HA
- More Cloud validations
- Majority Quorum (Potentially)

Linux | **Linux + AIX**

Pacemaker & Corosync   2/3

Which Pacemaker & Corosync version is supported ?
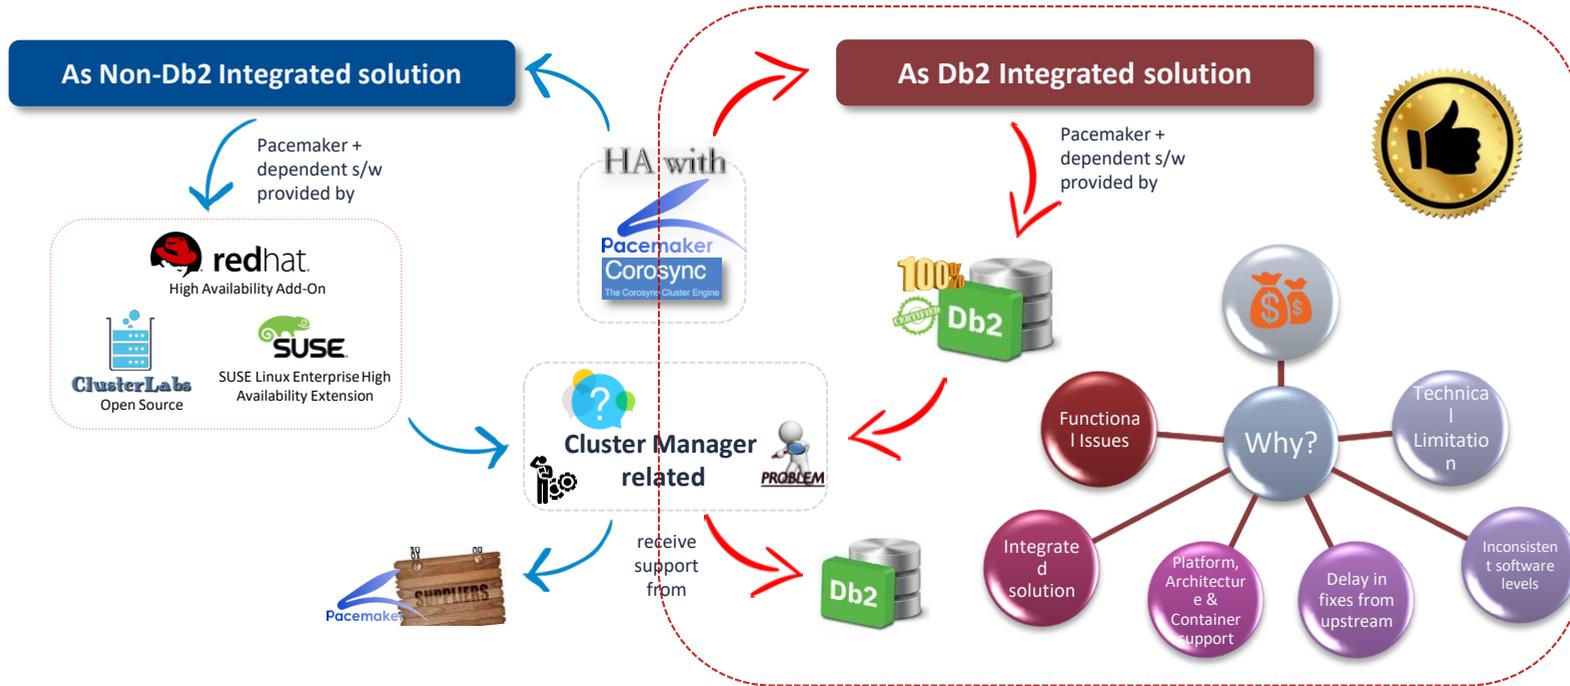
Pacemaker & Corosync   3/3

What platforms, Db2 Releases & deployment env are supported.

Architecture / Platforms / OS Version	TSA	Pacemaker
Intel / RHEL / 7.x	V11.5.4.0	No Plan
Intel / RHEL / 8.1	V11.5.4.0	V11.5.4.0
Intel / SLES / 12 SPx	V11.5.4.0	No Plan
Intel / SLES / 15 SPy	No	V11.5.4.0
Linux on IBM Z / RHEL 8.1	V11.5.4.0	V11.5.4.0
Linux on IBM Z / SLES 15 SP1	No	V11.5.4.0
POWER 8 RHEL 7.x	Yes	No Plan
POWER 8 & 9 / RHEL 8.1	No	~2021
POWER / AIX / 7.2 TL4	V11.5.4.0	~2021

| Environments	TSA	Pacemaker
On-premise DC	Yes	Yes
Non-containerized Private Cloud	No	Yes
Non-containerized Public Cloud	No	Yes, validated on AWS on Intel RHEL
Container	No	Not yet.

Note: For Pacemaker 11.5.4 is Tech Preview only, 11.5.5 will be 'GA' support
No plan to support Pacemaker with older version of RHEL (7.x) and SLES (12 SPx)
No plan to support Pacemaker earlier Db2 releases.

42

pureScale Concurrent Online Fix Pack / Mod Pack Update

IMPROVEMENT #1 NEW
Faster per host

Host 1   Host 2   · · ·   Host X

-33%   -33%   -33%

Time

Traditional serial online fixpack update but with each host taking less time

Host 1   Host 2   · · ·   Host X

Time

Traditional serial online fixpack update – host by host

IMPROVEMENT #2 NEW
Concurrent host online update while maintaining quorum

-33%   -33%

Host 1   Host X/2+1

Host X/2   Host X

Time

Concurrent online fixpack update

- Traditional serial host by host update benefit from improvement #1.
- Constant elapse time with concurrent update irrespective of cluster size (improvement #2)
- Go lightning fast with concurrent updates on half of the cluster at the same time!
- Improvement scales with increasing number of shared file systems!

Performance results done using a cluster with 36 file systems in a lab environment.

Periodic Automatic Cluster Topology Validation

Highlights:

- Speed boost in operational db2cluster commands *(Performance gain scales with number of shared file systems. e.g. with 36 shared FS, 33% gain was observed.)*

- Automatic periodic cluster topology validation by engine

Decouple topology validation from:

```
db2cluster -add -host ...
db2cluster -remove -host ...
db2cluster -cm -list -host ...
db2cluster -cfs -list -host ...
db2cluster -cm -set -tiebreaker ...
db2cluster -cfs -set -tiebreaker ...
db2cluster -cm -list -tiebreaker
db2cluster -cfs -list -tiebreaker
db2cluster -cm -enter -maintenance
db2cluster -cm -exit -maintenance
db2cluster -cfs -enter -maintenance
db2cluster -cm -exit -maintenance
...
```

validations move to

A new external "-topology" option in "db2cluster" to provide on-demand cluster topology validation

```
> db2cluster -verify -req -topology

--------------------------------------------------
Check 1: Verifying cluster topology. result: PASS
--------------------------------------------------
```

- New alerts will be generated or existing removed based on the results

Run as a new task in instance-wide "db2iperiodic" EDU.

- Executes every 24 hours (not-tunable dynamically)
- Only ONE member will execute at each 24-hour interval.
- New alerts will be generated or existing removed based on the results

44

Agenda

- A look back to Db2 11.5.0
- A look back at Db2 11.5.4
- A look back at Db2 11.5.5
- **A peek at some of the future 11.5.x+ candidates**

Performance – Columnar Trickle Insert   1/3

Goals:

Speed up "trickle" inserts and

reduce memory footprint.

Decrease size of very small tables.
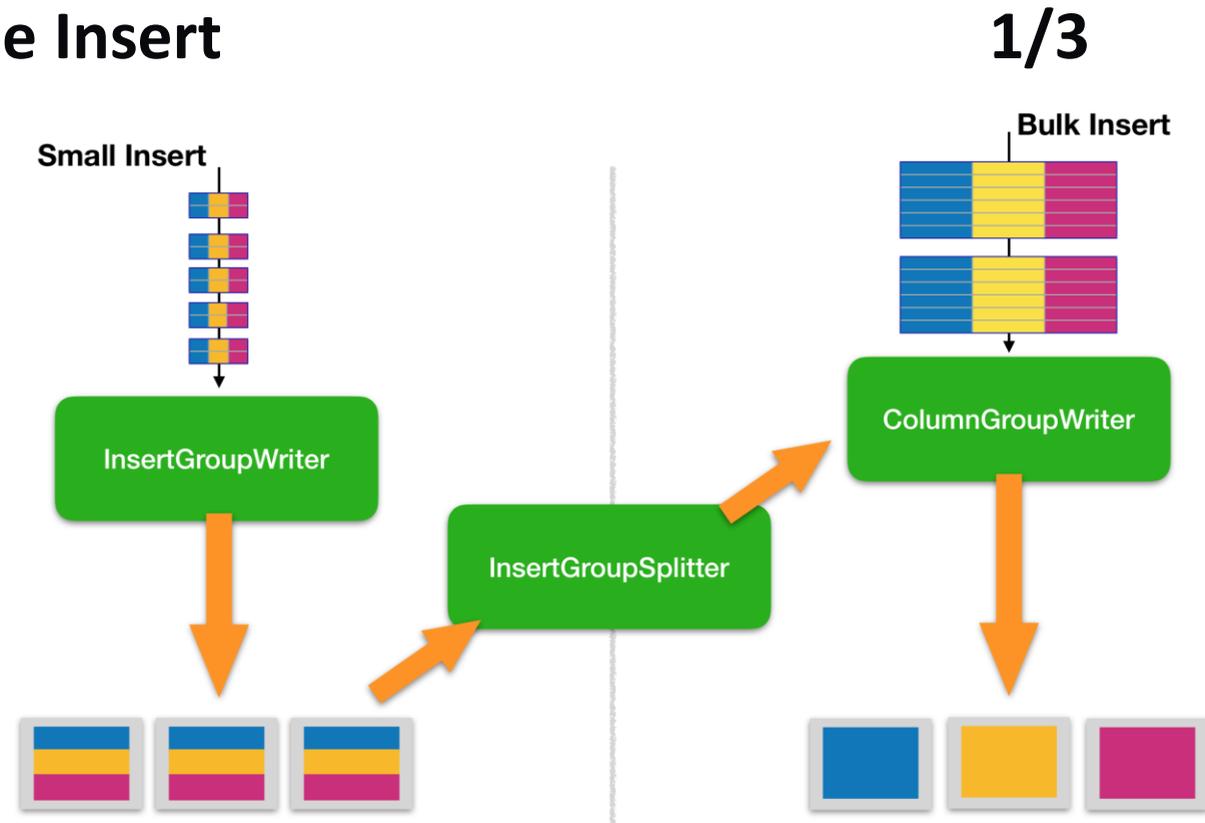
This will be off by default when it ships,   user will have to turn on as any table created or inserted into after it is enabled can no longer be accessed by previous level (no fallback support).

## Performance – Columnar Trickle Insert   2/3

Used only when small number of rows are being inserted (aka data trickling in).

Inserted rows are split to one or more "insert groups" – still columnar format just packed more columns per page.

Number of insert groups depends on types of columns, average length, etc. But generally an be much less than total number of columns.

Data going into these insert groups are not compressed.

Insert Groups

Insert Group Page 23; startTSN 42
- CG 0
- CG 1
- CG 2

Insert Group Page 111; startTSN 42
- CG 4
- CG 5
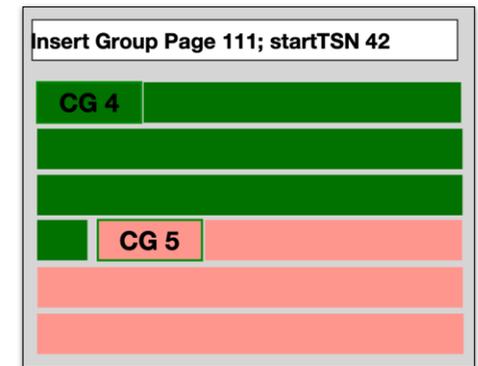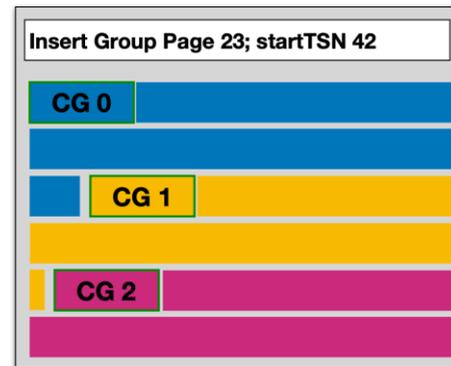
- Flexible assignments of column groups to insert groups.
  - Fixed-length vs. variable-length, large varchar columns.
  - Co-existence with other approaches, e.g., text compression.

## Performance – Columnar Trickle Insert   3/3

The insert group pages are almost always temporary -- a window of the most recent 'trickle' inserts.

Insert group pages/rows are automatically moved (aka split) to the single column per page format.

The split is done synchronously.   It's trigger as soon as it's predicted that full column group pages can be created.

Split Insert Group
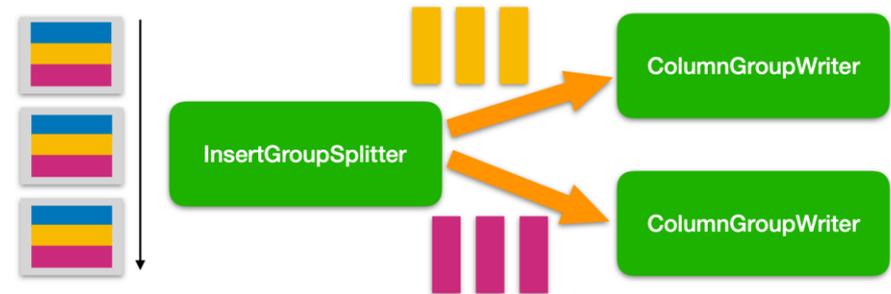
- **Synchronous** within a transaction (same approach as used for page compression).
- Tries to split when it thinks that a column group page can be filled.
  - Very **few pages** are **in Insert Group format**.
- Each column group is processed independently.

Other Columnar Candidates

System Maintained Temporal Support for Columnar Tables

Table automatic recompression to apply new compression features

- **Automatic analysis to determine when and what recompression to perform**
  - Currently have page-based string compression Type 1 and Type 2
  - Use statistics thresholds to automatically drive invocation
- **Recompress entire table (or subset of columns) to apply new page level compression features**

Real Time Statistics Collection

- **Full support for RTS adding sync statistics in addition to current fabrication support**

Query Performance and Memory Stability Enhancements.

- **Further Improve memory savings and performance for Join**
- **Aggregation distinct.**

Compatibility – Multi-tenancy

User-defined tenants

- A DBADM can create a tenant to set up an **independent catalog namespace** within a Db2 database

- All connections made to a database are initially associated with the SYSTEM tenant
  - A SET TENANT statement must be issued to associate a connection with a user-defined tenant

Initially targeted for Db2 Warehouse
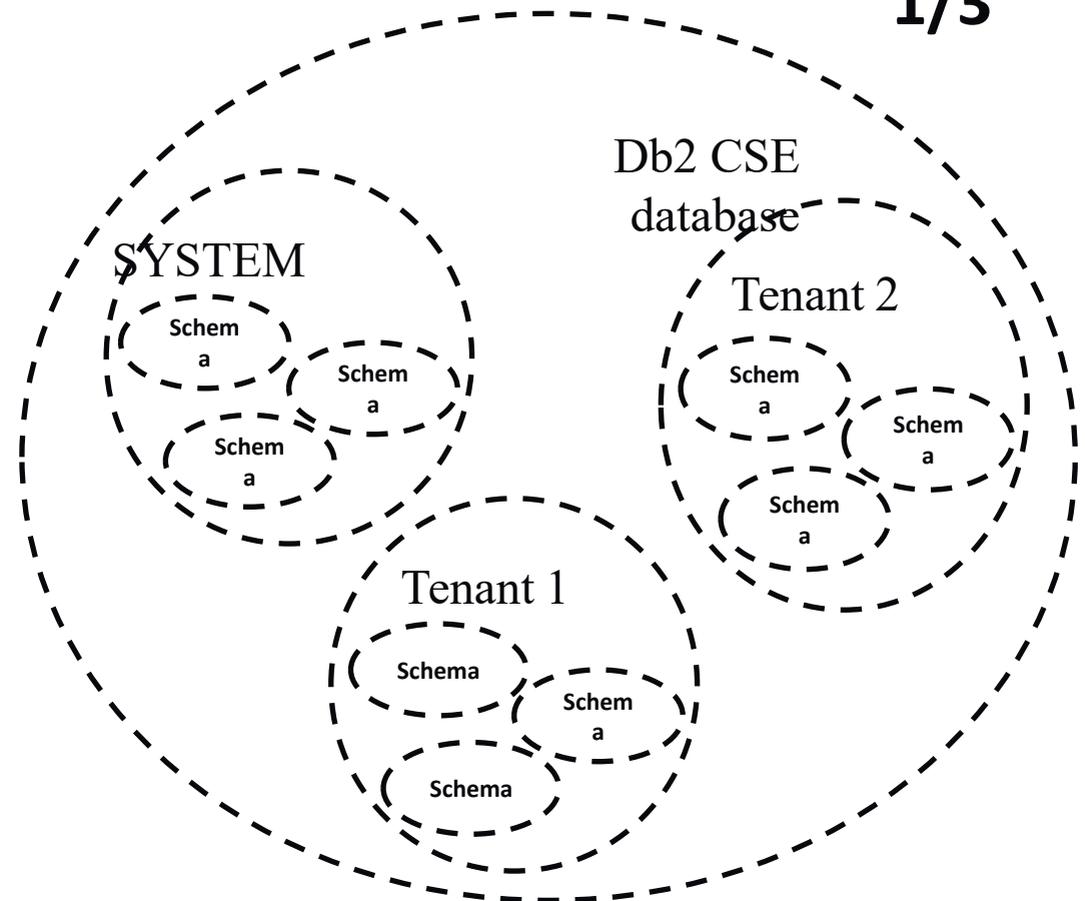
Compatibility – Multi-tenancy   **2/3**

A simple example (setup)

CREATE DATABASE TEST

CONNECT TO TEST

CREATE TENANT WORLD1
GRANT USAGE ON TENANT WORLD1 TO USER1

CREATE TENANT WORLD2
GRANT USAGE ON TENANT WORLD2 TO USER2

Compatibility – Multi-tenancy   3/3

A simple example

CONNECT TO TEST USING USER1   CONNECT TO TEST USING USER2

SET CURRENT TENANT = WORLD1   SET CATALOG WORLD2

CREATE TABLE MINE.T1 (C1 INT)   CREATE TABLE MINE.T1 (C1 CHAR(1))

INSERT INTO MINE.T1 VALUES (1)   INSERT INTO MINE.T1 VALUES ('A')

SELECT * FROM MINE.T1   SELECT * FROM MINE.T1
   -> 1   -> 'A'

Simplicity-- ALSM -- additional support

HADR Support

Extraction Logs for Online Backup, restore and rollforward

Log records from long
running transaction

TXID_EXTRACT.LOG

S0000000.LOG   S0000001.LOG   S0000002.LOG   S0000003.LOG

Available log files   Active log files
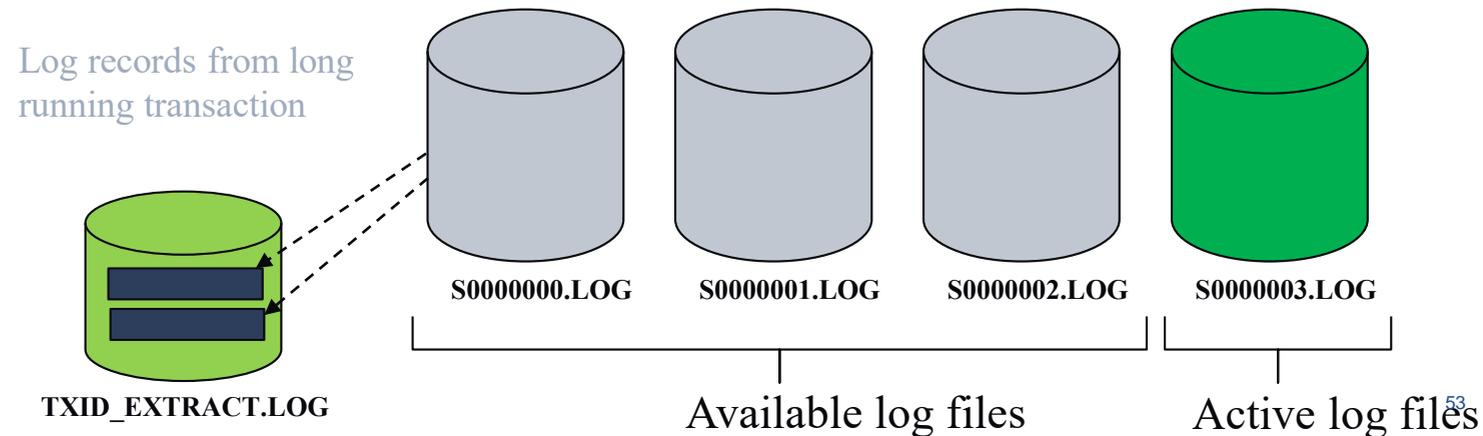
Availability Candidates

Next steps in non-pureScale HA with Pacemaker support

- Integrated bundling and install
- Active-Passive HA Configuration
- More cloud validations (IBM Cloud, Azure)
- DPF HA

NX842 hardware compression support with **encryption** for backup/restore.

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
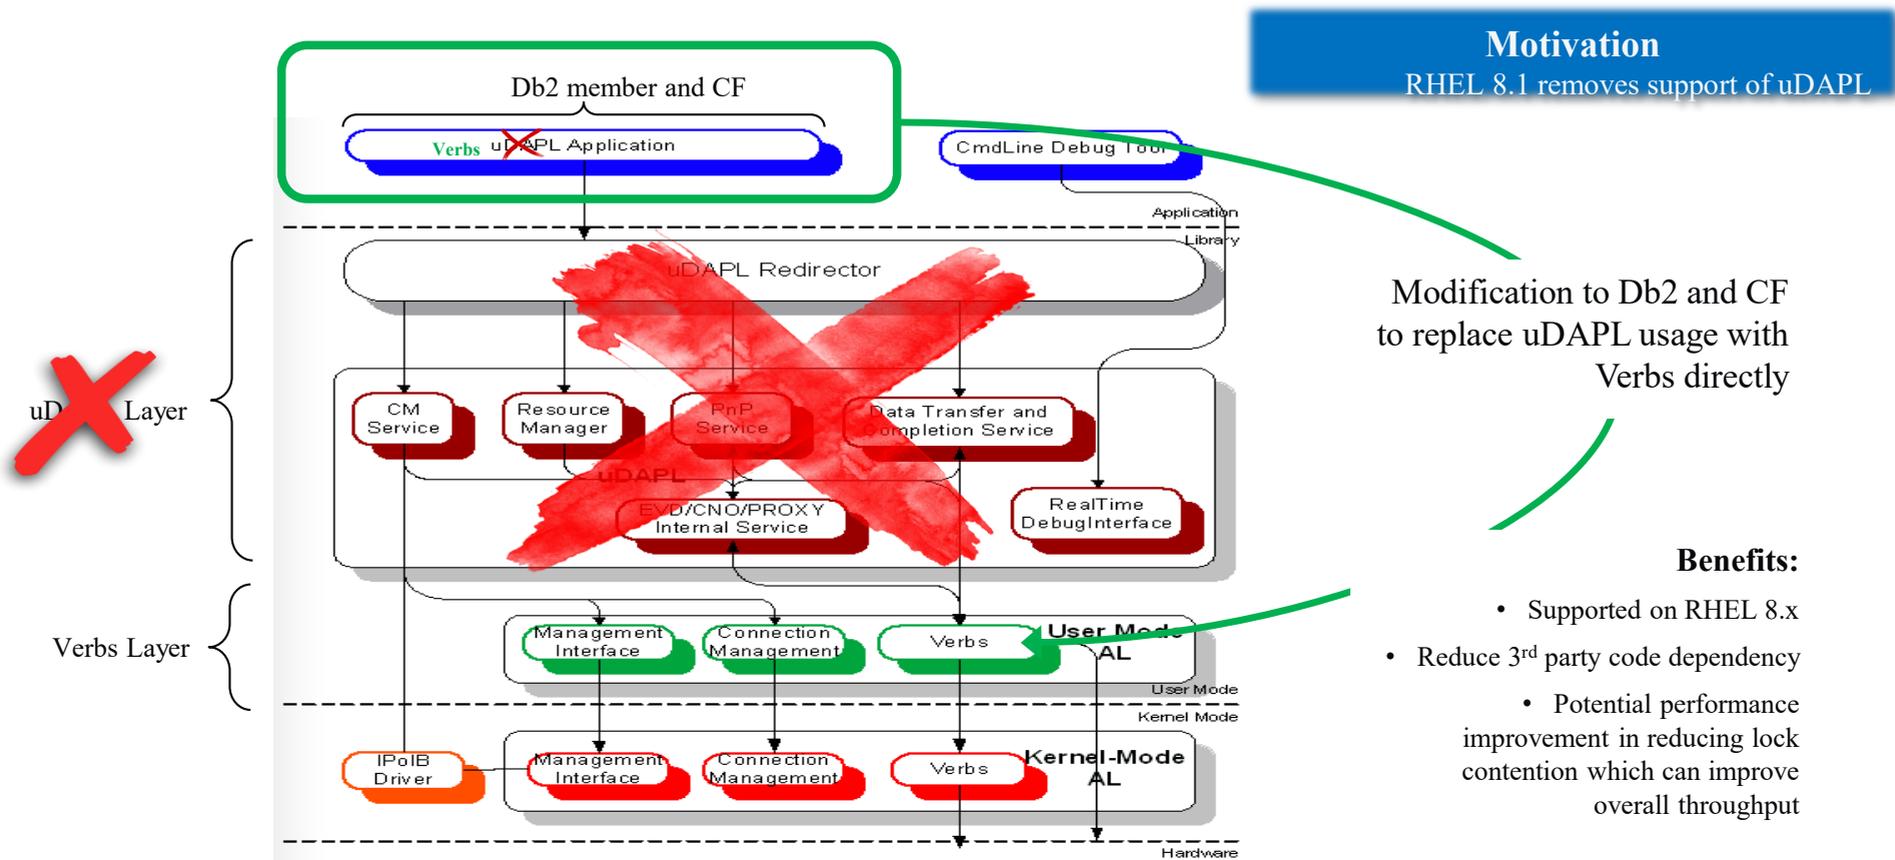
pureScale Candidates

Db2 pureScale on Linux on IBM Z – GA

Db2 "home-grown" RDMA adapter liveliness monitoring
- Removes one of the dependencies on RSCT – a step towards replacing TSA/RSCT with Pacemaker/Corosync in pureScale
- More accurate detection

RDMA Verbs Support -- RHEL 8.1 requirement for RDMA

SQL and Utility Enhancements

Improved Set Integrity Performance

• Additional parallelization of Set Integrity processing with Referential Integrity to speed up SI checking

Lifting SQL Limits Candidates

• Greater than 1012 columns per table (2048)
• Increase maximum total length of columns in a GROUP BY clause
• Increase Maximum total length of columns in an ORDER BY clause

Spatial Analytics Candidates

- GeoJSON Import and Export
- KML Import
- Box Filter support to improve performance
- ST_Ellipse support
- Spatial Aggregate Functions: Member and Union

IBM Db2 Resources

Information Resources:
- Db2 Roadmap - http://ibm.biz/AnalyticsRoadmaps
- Db2 RFE (Idea) Portal - http://ibm.biz/submitdb2idea
- Get latest Info- Subscribe to Db2 technical newsletter - http://ibm.biz/db2nlsignup
- Connect with the Db2 online community - http://ibm.biz/db2tribe
- Stay up to date with all the new demos on Youtube - http://ibm.biz/db2-youtube
- Interested in getting the latest beta version - http://ibm.biz/getdb2beta

Developer Resources:
- Db2 Developer Page to get started - http://ibm.biz/db2developer
- For Experienced Db2 developers, get your fav Db2 code sample on github - http://ibm.biz/db2github
- Want to try Machine Learning with Db2, check out - http://ibm.biz/learndb2
- Want details on Db2 Python Driver - http://ibm.biz/db2-drivers-python
- Want Details on Db2 PHP Driver - http://ibm.biz/db2-drivers-php
- Want Details on Db2 Node.js Driver - http://ibm.biz/db2-drivers-node
- Download the free Db2 python e-book - http://ibm.biz/db2pythonbook