



CCDUG



Db2 Latest from the Lab

Mike Springgay, IBM
Keri Romanufa, IBM

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.

AGENDA

2023 Deliverables

11.5.9 Overview

Deep Dives

Things starting with CS

Reference: 11.5.0 to 11.5.8 Overviews

AGENDA

2023 Deliverables

11.5.9 Overview

Deep Dives

Things starting with CS

Backup Deep Dives (skipped as covered in detail next 2 days)

Reference: 11.5.0 to 11.5.8 Overviews

Db2 2023 Roadmap

Q2 2023

- Db2 Warehouse Cloud Rack for P10

Q3 2023

- Db2 Warehouse Gen3 on AWS w/ **Db2 11.5.9**
- **Db2 11.5.9** container software technical preview (Db2 W only)
- Db2 Warehouse Cloud Rack for x86

Q4 2023

- **Db2 11.5.9** full software release (aka 11.5.9 mod pack)

AGENDA

2023 Deliverables

11.5.9 Overview

Deep Dives

Things starting with CS

Backup Deep Dives (skipped as covered in detail next 2 days)

Reference: 11.5.0 to 11.5.8 Overviews

Db2 11.5.9 Candidates



Hybrid & Multi-cloud

Containerization:

- Operator driven Db2 native backup and restore
- Db2 W Scaling - vertical by adjusting resources per pod
- Db2 W Scaling - horizontal by scaling out number of pods
- HADR role-aware Kubernetes service for Db2U. Route clients to primary.
- Support for Native Object Store – Db2 Warehouse only
- Support for ODF – Db2 Warehouse MPP only



Application Development

- Federation Enhancements
 - Support for AWS Aurora PostgreSQL
 - Support for PostgreSQL 12-15
 - Addition pushdown of:
 - timestamps to MongoDB
 - strings including trailing blanks comparison
 - numeric functions
 - date/time functions
 - CSE
 - Watsonx.data support via JDBC
- Client Enhancements:
 - Uplift of levels



Columnar Enhancements & Db2 Warehouse

- Support for tablespace storage to be on S3 (aka Remote Table spaces)
- Support for Open Data Formats: Iceberg, Parquet, AVRO, ORC
- Reference Architectures for Db2 Wh on Power and x86
- Online Synopsis Table Rebuild

* Delivered (1Q 2023) in latest containers

Db2 11.5.9 Candidates



Availability

- Ability to protect any History File entry from pruning
- Improved Replay Performance
- Db2 HADR with Pacemaker on Google Cloud Platform validated with SAP Workload



pureScale

- RHEL 8.8 support (might be post GA), SLES 15 FP5 support (might be post GA), AIX 7.3 TL1 support
- Z16 support
- CX-6 card support
- RoCE virtualization support on AIX 7.3
- CF reduction in latch contention with high read rates



Core Engine

- RHEL 9.2 support
 - RHEL 7 will no longer support integrated TSA (for HA or pS)
 - Ubuntu 22.4 support (no TSA/PCMK support)
- Advanced Data masking support
- Audit Logs to Object Store
- Restricted TCPIP Listener Mode
- ECMTS, Java, Gskit, TSA, GPFS and Pacemaker stacks updates

AGENDA

2023 Deliverables

11.5.9 Overview

Deep Dives

Things starting with CS

Backup Deep Dives (skipped as covered in detail next 2 days)

Reference: 11.5.0 to 11.5.8 Overviews

Db2 Warehouse – Remote Table Spaces

What: Use S3 (or compatible) storage for table space(s) storage

Why: #1 Reduce Storage Costs

- Object Storage is much much cheaper than block storage
- Data is still owned by Db2 and in Db2 optimized page formats
- A mix of block storage and object storage can still be used.
- Offset IO latency by use of NVME cache and optimizing IO patterns and storage layout

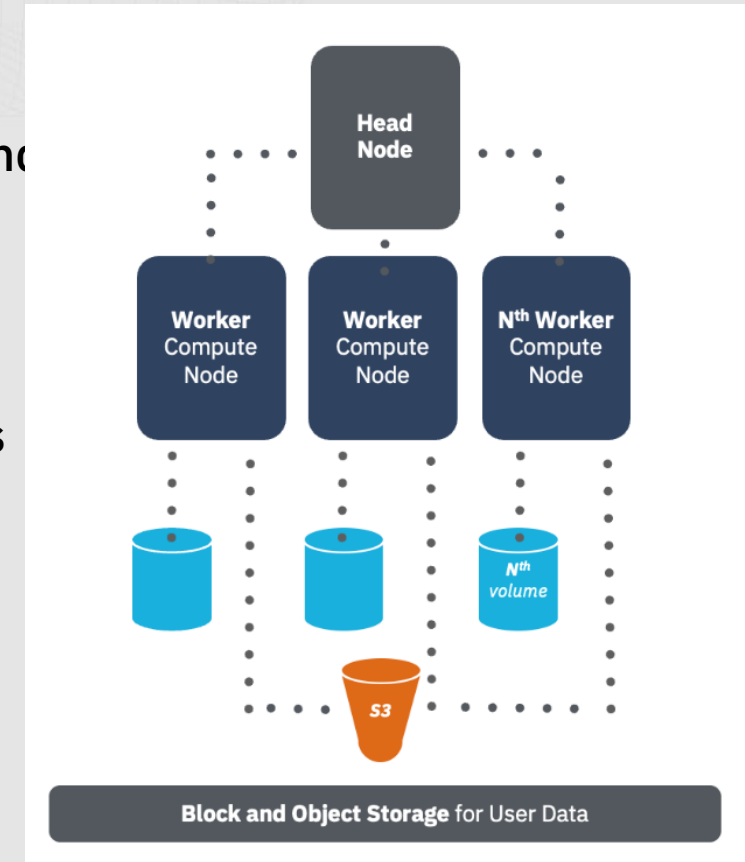
No applications nor workload changes needed

- simply create objects in appropriate (block or object store) table spaces

Initial Limitation:

Columnar tables only

Db2 Warehouse only (Db2U intel/amd and ppcle only initially)



DATALAKE Tables

1/2

A DATALAKE “Table” is a collection of files serialized using an **Open Data File (ODF)** format (TEXTFILE, ORC & Parquet) stored on remote storage (S3, COS). Data is not ‘owned’ by Db2.

The **metadata** of the table is stored in a metastore server:

- Location , Schema , Partition lists

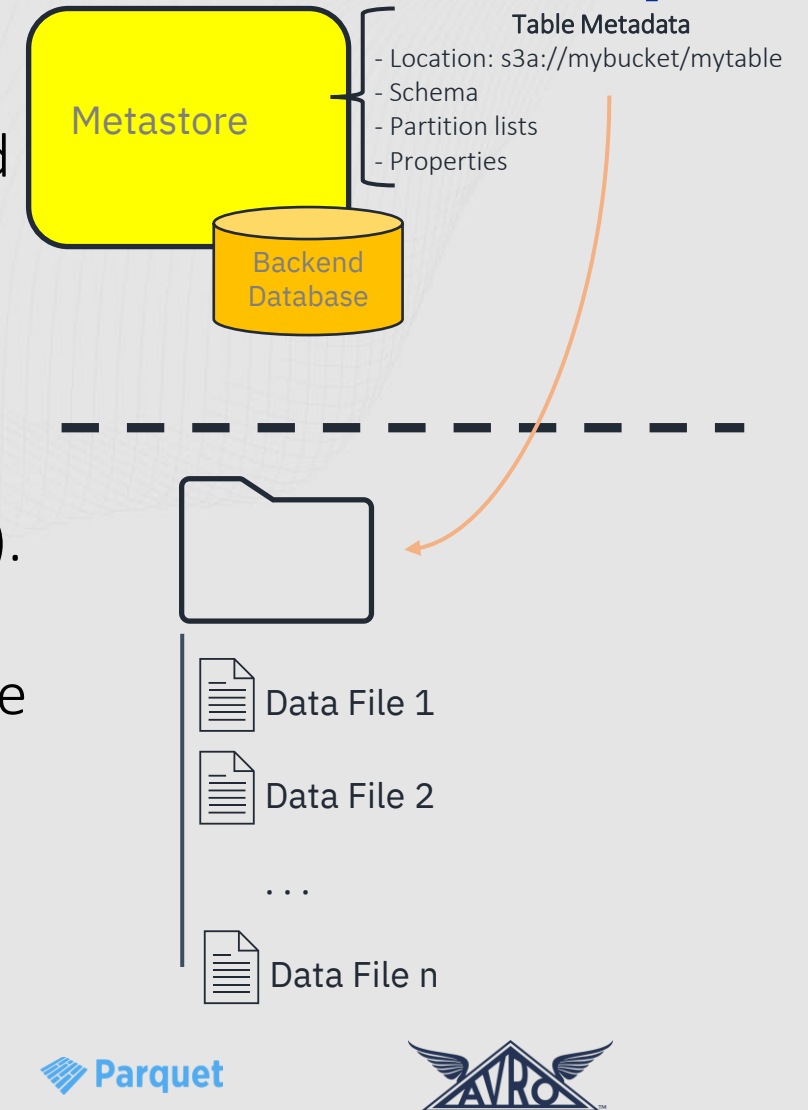
We will use Hive as the metastore . aka Hive metastore (HMS).

Db2 must query the metadata first and can proceed to read the data from remote storage

Benefits:

- Interoperability of open data formats
- Ease of use

Restrictions: Db2 Warehouse (MPP only / Db2U intel/amd and ppcle only initially)



DATALAKE Tables - Regular versus Iceberg

2/2

REGULAR -

Default type on CREATE DATALAKE TABLE

Does not support ACID properties for DML nor even simple transactions.

Recommended for read-only or append-only tables.

Formats: best performing (Parquet, ORC), TEXTFILE, AVRO, and JSONFILE

ICEBERG —

Specify STORED BY ICEBERG on the CREATE DATALAKE TABLE

Designed* to support ACID properties (does so via snapshots)

Recommended for scenarios which requires more flexibility on DML operations

Formats: Parquet, ORC , AVRO

*more details on limitations later

Note: Iceberg and regular DATALAKE Tables are additionally bound by Hive limits for things such as table, schema and column names and other limits. In addition, see docs for mappings and limits when it comes to data types when creating Regular and Iceberg DATALAKE tables..

Protect History Entries from Auto Deletion / Pruning

Prior to 11.5.9 only history file entries related to recovery can have their status marked with "X" (DB2HISTORY_STATUS_DO_NOT_DELETE) to protect them from pruning.

As of 11.5.9, all history entry types can now have their status set to X.

Hint: use carefully, as using on too many entries can result in unexpected growth of the history file.

Restricted TCP listener mode

1/3

Purpose:

- Allow secured application connections using connect type 2

• Problem:

- TCP listener is not running (SVCENAME, DB2COMM)
- Applications use connect type 2 using secure port only
- Sync Point Manager requires unsecured TCPIP (DRDA protocol)
- Distributed transaction processing is not possible and db2diag.log message is logged

Restricted TCP listener mode

2/3

Solution:

- TCPIP listener can be started in restricted mode, dropping any connection request other than:
 - resynchronization (SPM)
 - DRDA ping

How to Setup:

- configure SVCENAME in DBM CFG
- DB2COMM registry variable does **not** contain value TCPIP

Logging:

- One EVENT db2diag.log message in instance start (if configured)
- One INFO db2diag.log if clients attempt to connect and are rejected logging client IP

Restricted TCP listener mode

3/3

Example of rejected application connect on client:

```
SQL30081N A communication error has been detected. Communication
protocol being used: "TCP/IP". Communication API being used:
"SOCKETS". Location where the error was detected: "<IP
address>". Communication function detecting the error: "recv".
Protocol specific error code(s): "*", "*", "0". SQLSTATE=08001
```

Online Columnar Synopsis Table Rebuild

1/2

- Purpose of rebuilding synopsis table:
 - Reduce synopsis table excessive storage usage due to sparsity in data pages
 - Recovery from issues with synopsis table if encountered
- Restriction:
 - When rebuilding a synopsis table, the corresponding base table is in EXCLUSIVE mode, not available for read or write. (So database ONLINE, but table not accessible till complete.)

Online Columnar Synopsis Table Rebuild

2/2

Rebuild synopsis table syntax :

```
REORG SYNOPSIS FOR TABLE <base table> REBUILD
```

Example:

```
REORG SYNOPSIS FOR TABLE customer REBUILD
```

The REORG SYNOPSIS command needs to be the first statement in a UOW. If not, SQL0428N will return.

If the REORG SYNOPSIS command does not complete successfully (SQL2223N), the synopsis table will remain in an unusable state.

- The synopsis table will not be included in any query plan
- Querying the synopsis table will receive an error (SQL0668N reason code 12).

Data Masking

1/3

Purpose: provide enhanced data privacy capabilities

Adds built-in masking function that support:

- Redaction
- Partial Redaction
- Substitution
- Obfuscation

Can be combined with Row and Column Access Control to provide advanced data privacy

```
DATA_MASK (expression, mask-type,  
            mask-parameters, mask-format, seed)
```

Expression

- Expression that returns value to be masked

Mask-type

- An expression that specifies masking operation to be perform

Mask-parameters

- String that provides additional parameters to influence masking operation

Mask-format

- What format to apply when Obuscation format preserving / format preserving fabrication

Seed

- A string value to use as seed to hashing function

Data Masking

3/3

MaskType	Masking Name	Data-Type Support	Description
0	REDACT	CHAR , VARCHAR*	redact string via mask-parameters
1	REDACT PRESERVE LENGTH	CHAR , VARCHAR*	redact character via mask-parameter
2	SUBSTITUE	All support data types	Strings SHA-256 hash & base64 encode. Others within type range
3	PARTIAL REDACT	CHAR , VARCHAR*	Pattern via mask-parameters
4	FORMAT PRESERVING	CHAR , VARCHAR, DATE TIMESTAMP*	matches format via mask-format
5	DATE AGING	DATE and TIMESTAMP*	mask-parameters how date is aged
7	IDENTIFIER	CHAR , VARCHAR*	Alphabets and digits masked other characters remain as-is
9	FORMAT PRESERVING PARTIAL	CHAR , VARCHAR*	Only for Email currently
10	FORMAT PRESERVING FABRICATION	CHAR , VARCHAR, DATE TIMESTAMP*	obfuscation without input validation
11	NUMERIC SHIFT	Numeric Types	mask-parameters integer percentage to shift input value by
* All other data-types redacted to default corresponding value			

AGENDA

2023 Deliverables

11.5.9 Overview

Deep Dives

Things starting with CS

Backup Deep Dives (skipped as covered in detail next 2 days)

Reference: 11.5.0 to 11.5.8 Overviews

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, we've tried 2 options:

- Fixpack - increment just the FP level
 - only done once in life of 11.5 (11.5.5.1)

(same as 11.1)

- CSB, no numbering change, it is a cumulative SB level
 - Maintained continually on 11.5.4 and later

Db2 11.5: CSB vs past “special build”

How are they alike:

- both are how customers get APARs fixes “asap”
- both are identified by a special build ID (versus change in mod or FP level)

How do they differ (aka why is the CSB much better).

- C stands for Cumulative / Continuous
 - The CSB is all APARs (including security) fixed since the last mod pack.
- Any/all customers pick up the same stream, the previous special build model resulted in “1-of” builds for customers
- Security SB is just a CSB w/ the latest security fixes-- they are part of the cumulative stream.
- Being a continuous stream, it gets continuous testing including pipeline, functional, performance and stress; previous SB had minimal testing (decided per fix).

If a customer is installing a new Db2 instance or moving up to 11.5; it is recommended that they pick up the latest CSB stream (will include the latest).

Note we still have ability to do 1-ofs (and they are call unofficial SB) but used almost exclusively for debug code or partial fixes not meant or ready for everyone’s use.

Db2 11.5: going forward

1/2

Rely on CSB stream, no fix packs.

Fix packs simply package many APARs together.

Good:

- as a periodic encouragement to move up; “look here’s a bunch of great fixes”
- to give a target level/date to help support/justify a plan to update (to mmgt, other teams, etc)

But:

- often dates don’t align with when customers want to do a new install, or upgrade or even pre-planned updates.
- it is stale the moment it ships, let alone by the time the next one comes out
 - e.g mod 11.5.8 released it’s first CSB with ~50 additional APAR fixes the same day
- same issue tends to be hit multiple times till next FP out (which can be months)

Db2 11.5: going forward

2/2

Being able to pick-up the latest CSB solves most of the issues.

Picking a specific CSB level at some interval (current feedback is in range of ~3 months) with priority on security ones to document as “suggested” levels and push customer to move to them similar as they would fix packs, they can of course still pick to go to latest level by the time they implement.

This is still being rolled out! We do have 1 more mod pack coming now!

KIs need to be updated to include CSB level it's fixed in.

~~—Ensure KIs are published and listed to match CSBs available~~ **DONE ***

~~—Ensure there is a document similar to the security vulnerabilities page to
——point to the latest and ‘selected’ CSB levels.~~ **DONE ***

Some documentation that is fixpack-based needs updating.

* see <https://www.ibm.com/support/pages/node/6856211>

Announcement: New extended support

A new Extended Support was announced for all of IBM Software in April:

<https://www.ibm.com/support/pages/node/7008589>

Highlights: 4 years of extended support.

First year includes defect fixes and security patches

Remaining 3 years are “usage and existing fixes” only

Shortly after, Db2 announced end of Continuing Support as of April 30, 2024, and that all levels that were still under CS (9.7 and up) would also have the new ES to be available April 30, 2025 (for 11.1 2026). Note for this its: usage / existing fixes only plus limited to Db2-only not any embedded products (similar limit as to CS)).

For more details see:

<https://www.ibm.com/support/pages/db2-distributed-end-support-eos-dates>

<https://www.ibm.com/support/pages/node/7008589>

AGENDA

2023 Deliverables

11.5.9 Overview

Deep Dives

Things starting with CS

Backup Deep Dives (skipped as covered in detail next 2 days)

Reference: 11.5.0 to 11.5.8 Overviews

AGENDA

Backup Deep Dives

- Remote Table spaces
(aka columnar tables on cloud object store)

Db2 Warehouse – Remote Table Spaces

What: Use S3 (or compatible) storage for table space(s) storage

Why: #1 Reduce Storage Costs

- Object Storage is much much cheaper than block storage
- Data is still owned by Db2 and in Db2 optimized page formats
- A mix of block storage and object storage can still be used.
- Offset IO latency by use of NVME cache and optimizing IO patterns and storage layout


No applications nor workload changes needed

- simply create objects in appropriate (block or object store) table spaces

Initial Limitation:

Columnar tables only

Db2 Warehouse only (Db2U intel/amd and ppcle only initially)



Db2 Warehouse Storage Architecture



Db2 Warehouse Storage Architecture



High Performance



Elastic



Monolithic Storage




Expensive



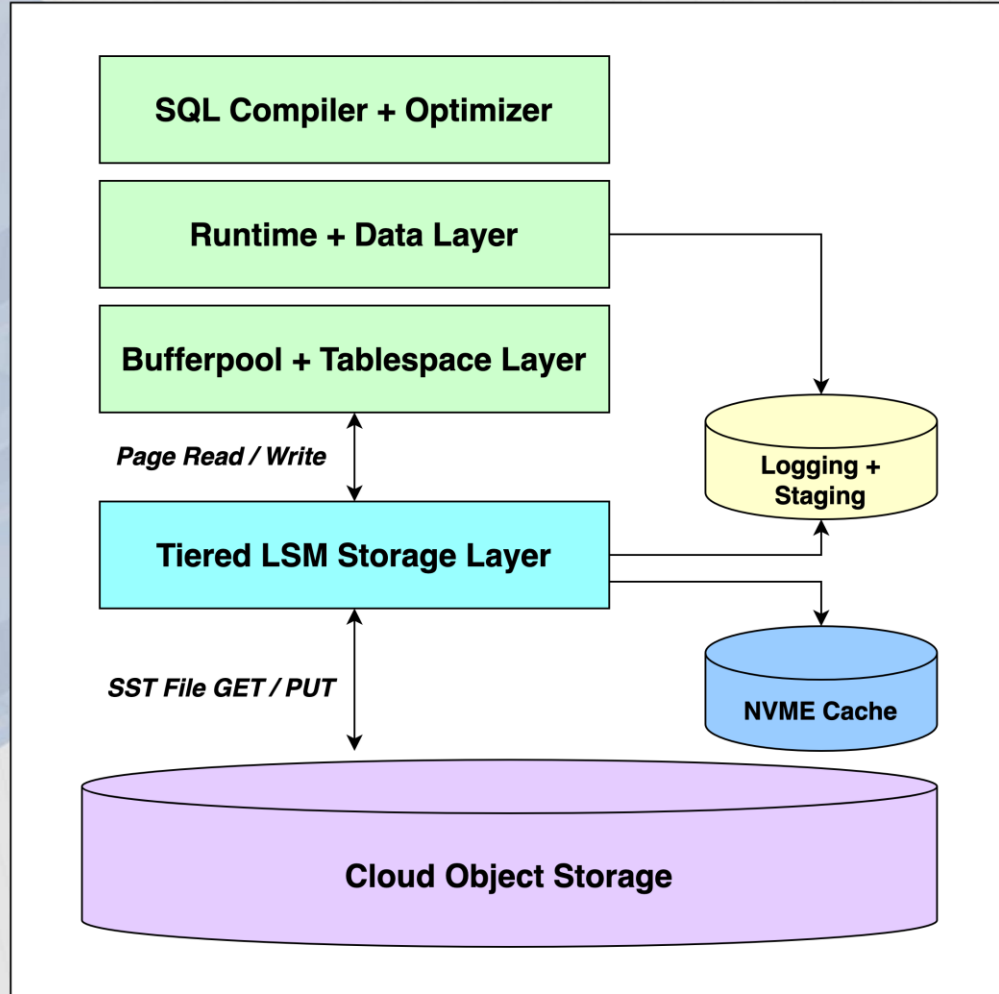
Db2 Warehouse Storage Architecture

with Cloud Object Store Table Spaces

- ✓ ***High Performance***
- ✓ ***Elastic***
- ✓ ***Decoupled Storage***
- ✓ ***Cost Efficient***




Architectural Changes for Remote Table Space



- Existing Db2 component stack down through buffer pool + table space layer
- Existing Db2 logging maintains high performance for trickle feed
- Three new elements in new native cloud storage layer:
 1. An LSM tree storage organization to efficiently store Db2 native pages on cloud object storage.
 2. A novel data clustering technique that exploits the self-clustering capabilities of the LSM tree.
 3. A multi-tiered cache that adds a local NVMe component to enable high performance query processing and bulk ingest.


Remote Tablespace - Optimizations for Performance

Today



1000 page writes to fast
network attached block
storage


11.5.9



1000 pages writes to local NVME
(ephemeral)
+ 1 PUT to S3

BDI Mixed Workload – 10 TB – Cold Start

Speedup WHoC Gen 2 vs. WHoC Gen3 (w/ Remote Table Spaces)



* Customer results may vary. These results are based on internal tests with pre-production level of code. The tests were set up on two equivalent environments with 24 database partitions on two AWS EC2 nodes, each with 48 cores, 768 GB memory and a 25 Gbps network interface. In the case of the Db2 Warehouse Gen3 environment, this adds 4 NVMe drives per node for a total of 3.6 TB, with 60% allocated to the on-disk cache (180 GB per database partition, or 2.16TB total).

Catalog / monitor - Remote Table Spaces

Table functions and catalogs show remote object store & NVMe Caching usage.

```
SELECT VARCHAR(TBSP_NAME, 30) AS TBSP_NAME, MEMBER, TBSP_TYPE, CACHING_TIER
FROM TABLE(MON_GET_TABLESPACE('',-2)) AS T
```

TBSP_NAME	MEMBER	TBSP_TYPE	CACHING_TIER
OBJSTORESPACE1	0	DMS	ENABLED
OBJSTORESPACEUTMP1	0	DMS	ENABLED

```
SELECT VARCHAR(CONTAINER_NAME,40) AS CONTAINER_NAME, VARCHAR(TBSP_NAME,20) AS TBSP_NAME,
CONTAINER_TYPE FROM TABLE(MON_GET_CONTAINER('',0))
```

CONTAINER_NAME	TBSP_NAME	CONTAINER_TYPE
DB2REMOTE://IBMDEFAULTREMAlias	OBJSTORESPACE1	DB2REMOTE
DB2REMOTE://IBMDEFAULTREMAlias	OBJSTORESPACEUTMP1	DB2REMOTE

* Note: these match the table spaces create by DEFAULT in WoC Gen 3

Remote Table Space – 11.5.9 initial Restrictions

- System temporary tablespaces are not supported in remote storage.
- Row organized tables cannot be created in remote table spaces.
- Indexes for column-organized tables cannot be created in remote table spaces (use the INDEX IN clause to define a non-remote table space in order to create an index). This also applies to enforced key constraints.

WoC Gen 3: There are 2 automatically created table spaces (OBJSTORESPACE1,OBJSTORESPACEUTMP1) No create/drop allowed. Same for the remote access alias (IBMDEFAULTALIAS)

Db2 W : limits on number aliases and/or table spaces and creation/dropping of remote table spaces still **TBD**.

AGENDA

Backup Deep Dives

- DATALAKE tables

DATALAKE Tables

A DATALAKE “Table” is a collection of files serialized using an **Open Data File (ODF)** format (TEXTFILE, ORC & Parquet) stored on remote storage (S3, COS). Data is not ‘owned’ by Db2.

The **metadata** of the table is stored in a metastore server:

- Location , Schema , Partition lists

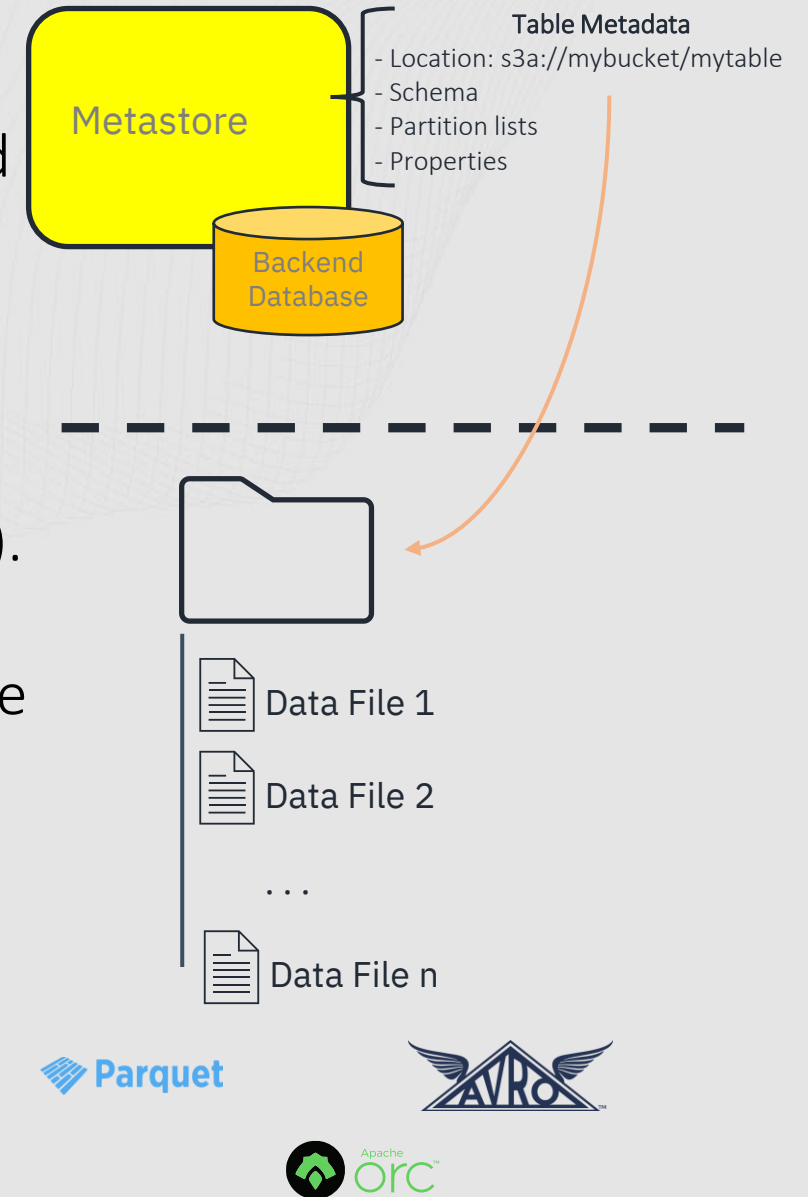
We will use Hive as the metastore . aka Hive metastore (HMS).

Db2 must query the metadata first and can proceed to read the data from remote storage

Benefits:

- Interoperability of open data formats
- Ease of use

Restrictions: Db2 Warehouse (MPP only / Db2U intel/amd and ppcle only initially)



DATALAKE Tables - Regular versus Iceberg

REGULAR -

Default type on CREATE DATALAKE TABLE

Does not support ACID properties for DML nor even simple transactions.

Recommended for read-only or append-only tables.

Formats: best performing (Parquet, ORC), AVRO, JSONFILE and TEXTFILE

ICEBERG —

Specify STORED BY ICEBERG on the CREATE DATALAKE TABLE

Designed* to support ACID properties (does so via snapshots)

Recommended for scenarios which requires more flexibility on DML operations

Formats: Parquet, ORC , AVRO

*more details on limitations later

Note: Iceberg and regular DATALAKE Tables are additionally bound by Hive limits for things such as table, schema and column names and other limits. In addition, see docs for mappings and limits when it comes to data types when creating Regular and Iceberg DATALAKE tables..

Regular DATALAKE Table DDL Samples

CREATE - This creates the table in the Db2 catalogs and the HMS. If the data files already exist, pre-existing data can be read. Whether the files pre-exist or not, new data can be appended. See docs for additional limits/features, including partitioning options and datatype support.

```
CREATE DATALAKE TABLE my_datalake_table(id INT, name
VARCHAR(8)) STORED AS PARQUET LOCATION 'DB2REMOTE://mybucket-
alias//my_datalake_table'
```

DROP - This removes the definition from the db2 catalogs and HMS. See the docs for details on how to control whether the data files are dropped as part of the DROP command or not.

```
DROP DATALAKE TABLE my_datalake_table
```

Some **ALTER** operations are also supported.

The CREATE/DROP/ALTER DDL are outside the scope of the transaction and rollback/commit at statement level. For CREATE TABLE AS it acts as 2 individually committed statements (CREATE+INSERT)

Regular DATALAKE Use Cases

First step is to create the DATALAKE Table which includes ensuring access to the S3 bucket & using a stored procedure to set up the storage access alias. Once created there are 3 common use cases.

Queries with DATALAKE Table(s)

DATALAKE tables can be used in Db2 queries like (and along with) Db2 tables.

Loading data from DATALAKE Tables into Db2 tables.

Data can loaded into Db2 tables from DATALAKE tables by an INSERT into the Db2 table with the SELECT FROM clause accessing a DATALAKE table.

Exporting data from Db2 into DATALAKE Tables

Data can put into DATALAKE Tables by using the INSERT statement into the DATALAKE table with the SELECT FROM clause accessing a Db2 Table.

Limitations

Anything that would be to related ownership or responding to changing data: CREATE INDEX, REORG, TRIGGERS, ..
No transaction or ACID support

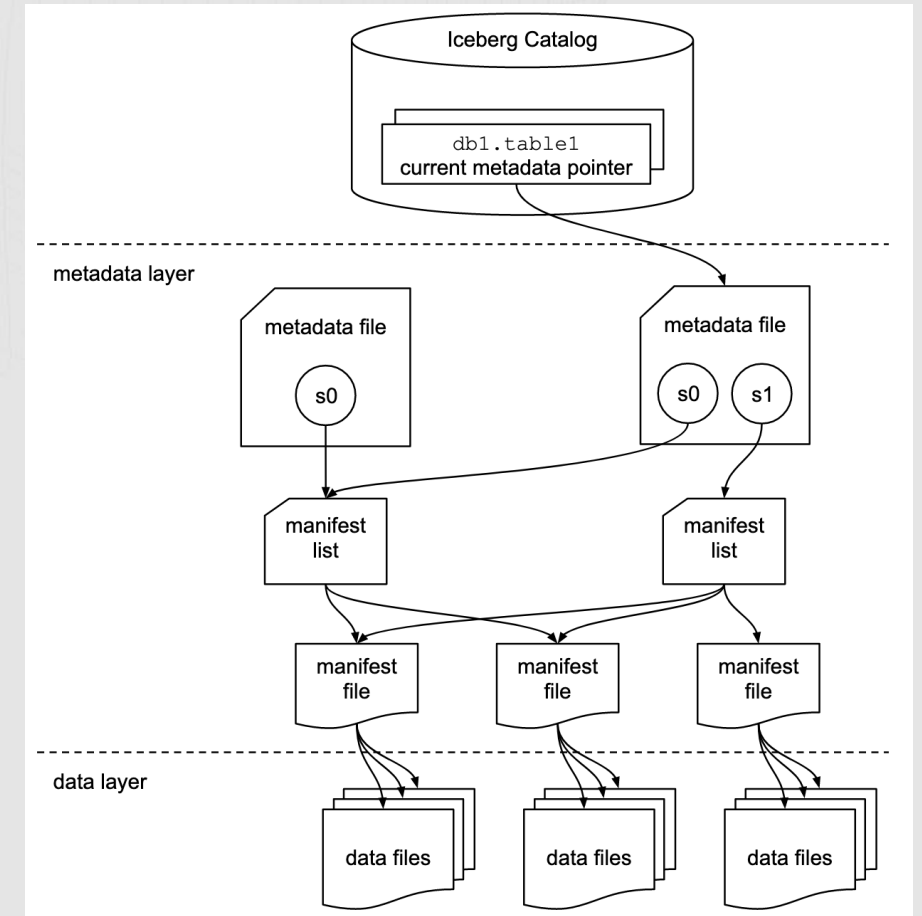
Apache Iceberg

An Open Data Table format



Full **open-source, Open Data Table format**, quickly becoming an **industry standard**

Relies on Open Data File formats for storage, but provides an additional layer of metadata for enhanced capabilities



Iceberg Support in Db2 Warehouse

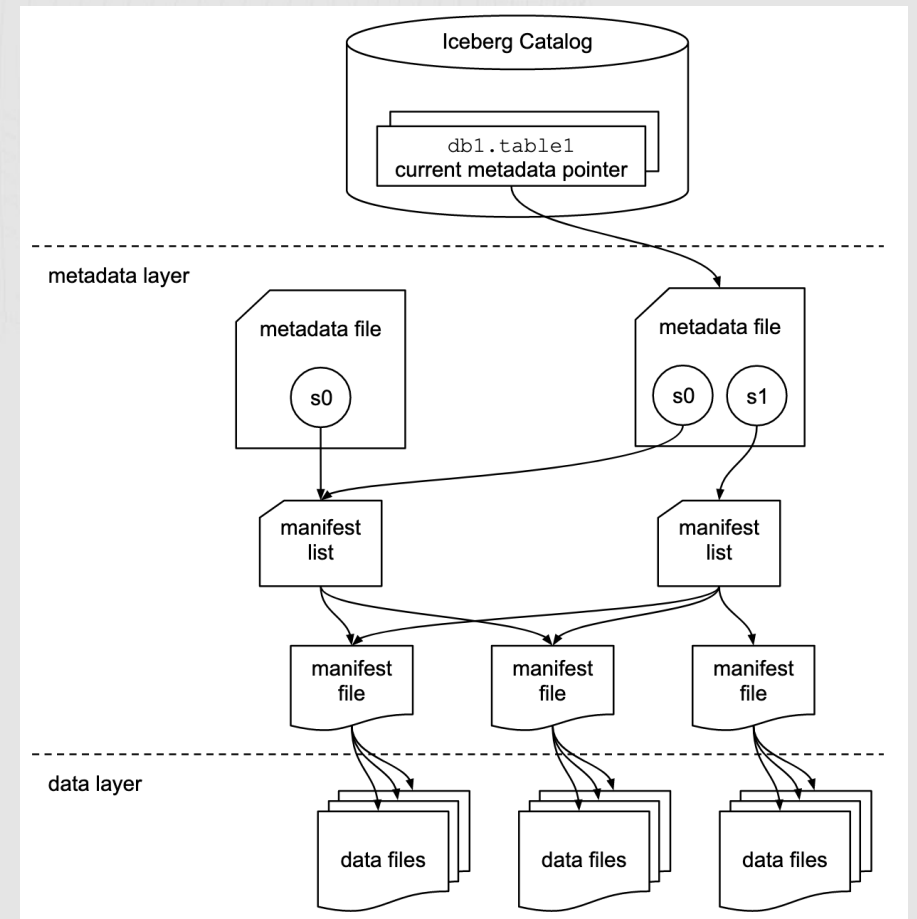


Support for:

- CREATE (DDL or GUI)
 - Includes partitioning support
- SELECT (aka READ)
- INSERT
- Snapshot Rollback
- Most schema evolution

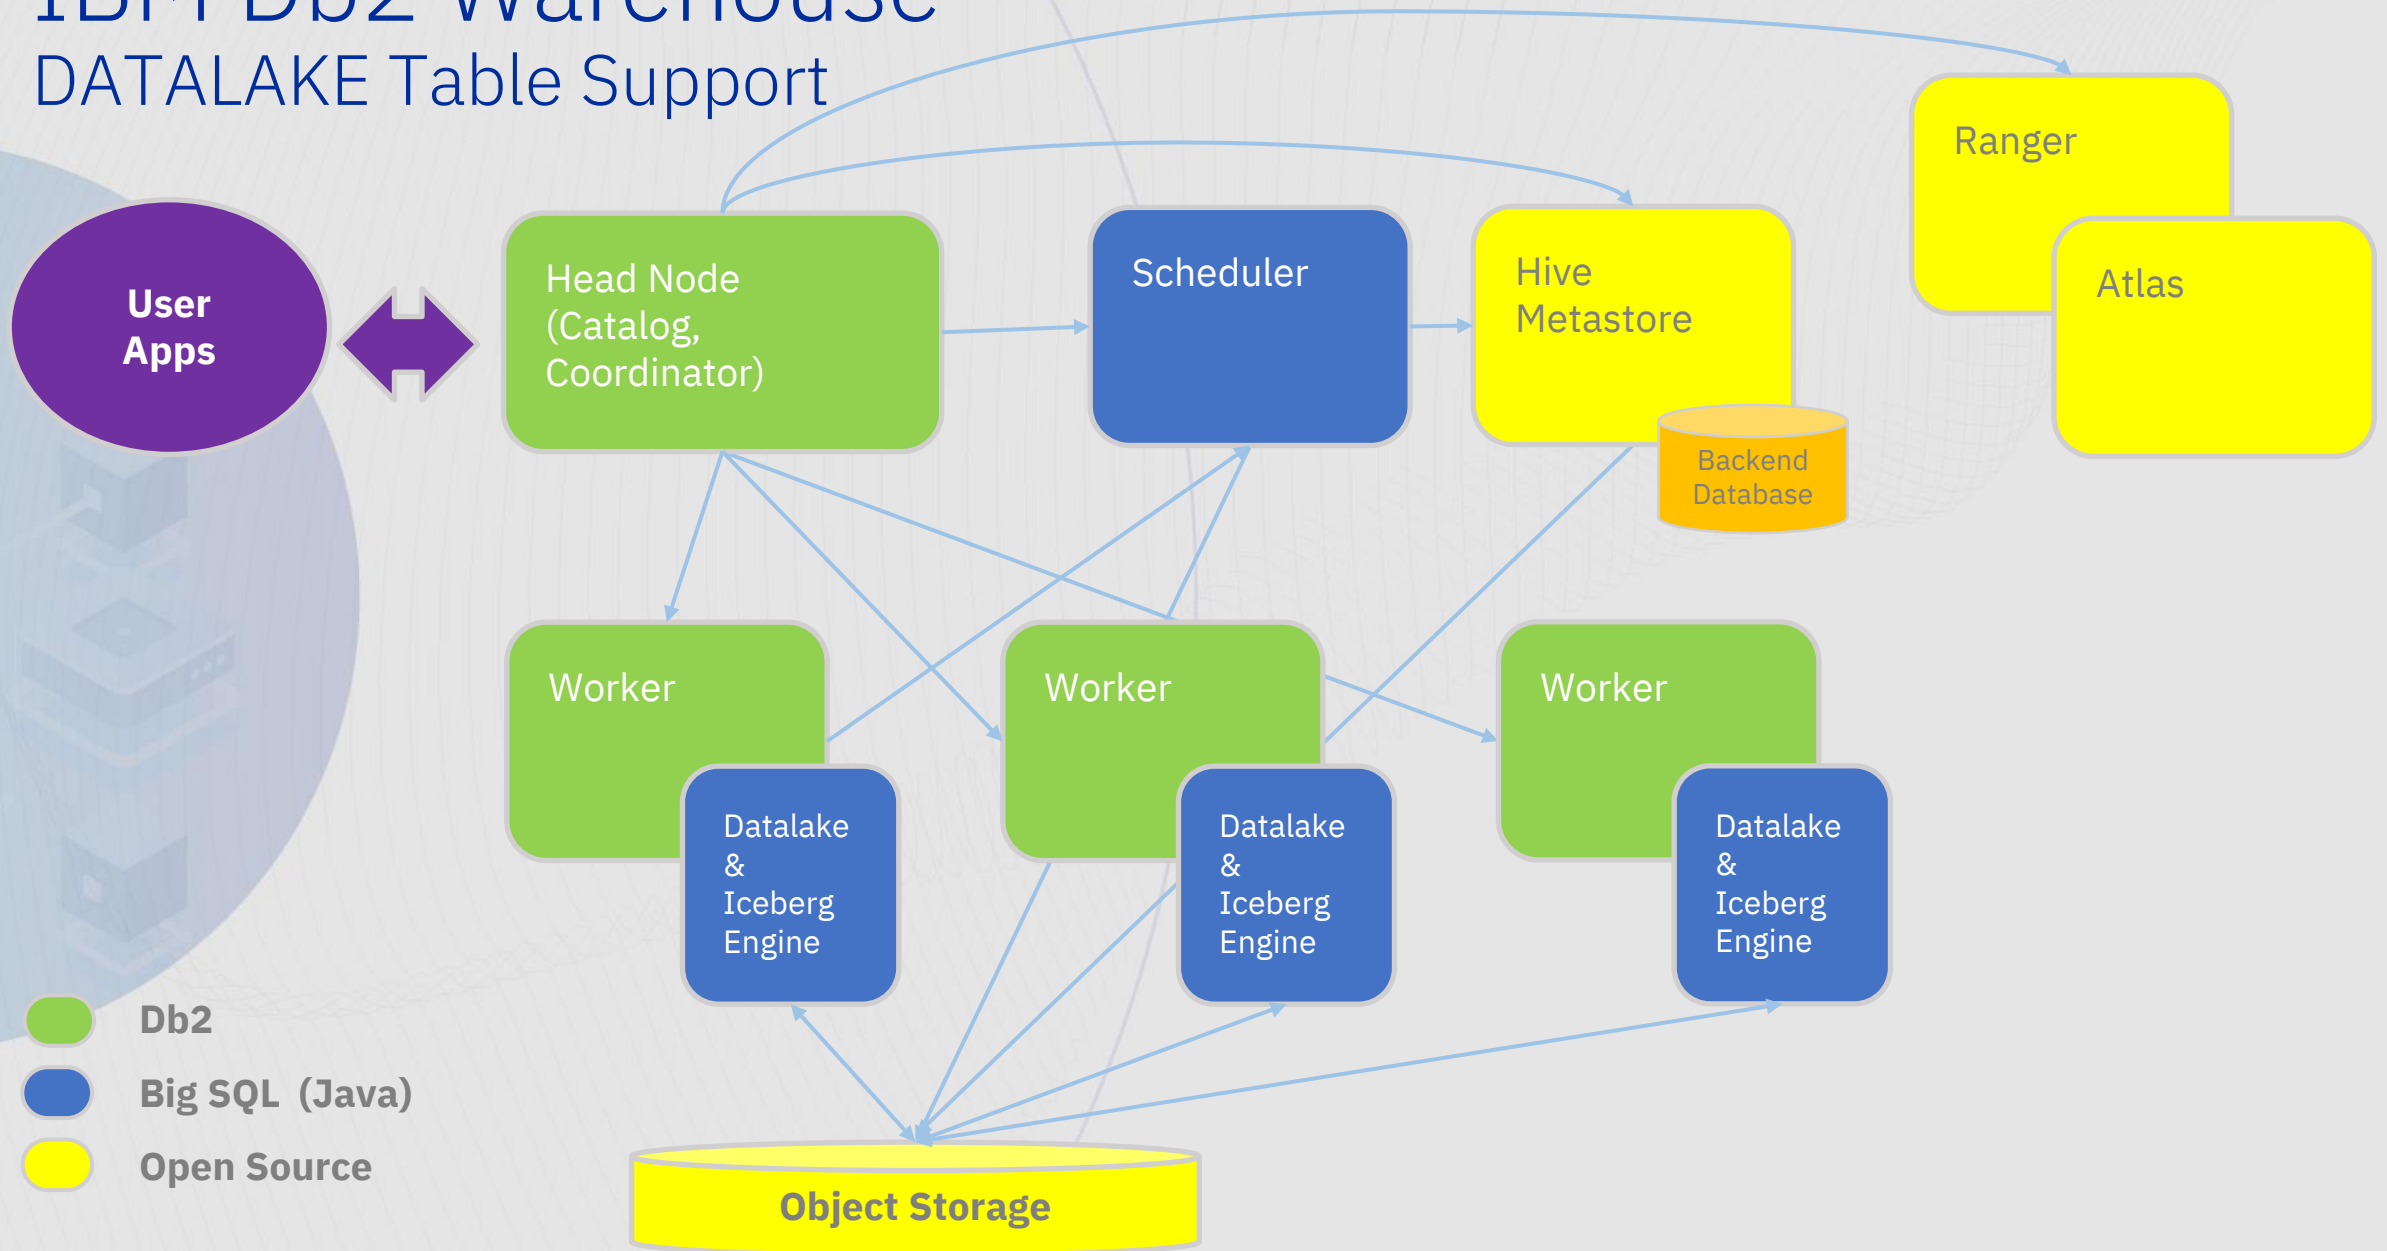
Current Restrictions:

- No UPDATE / DELETE
- No time travel



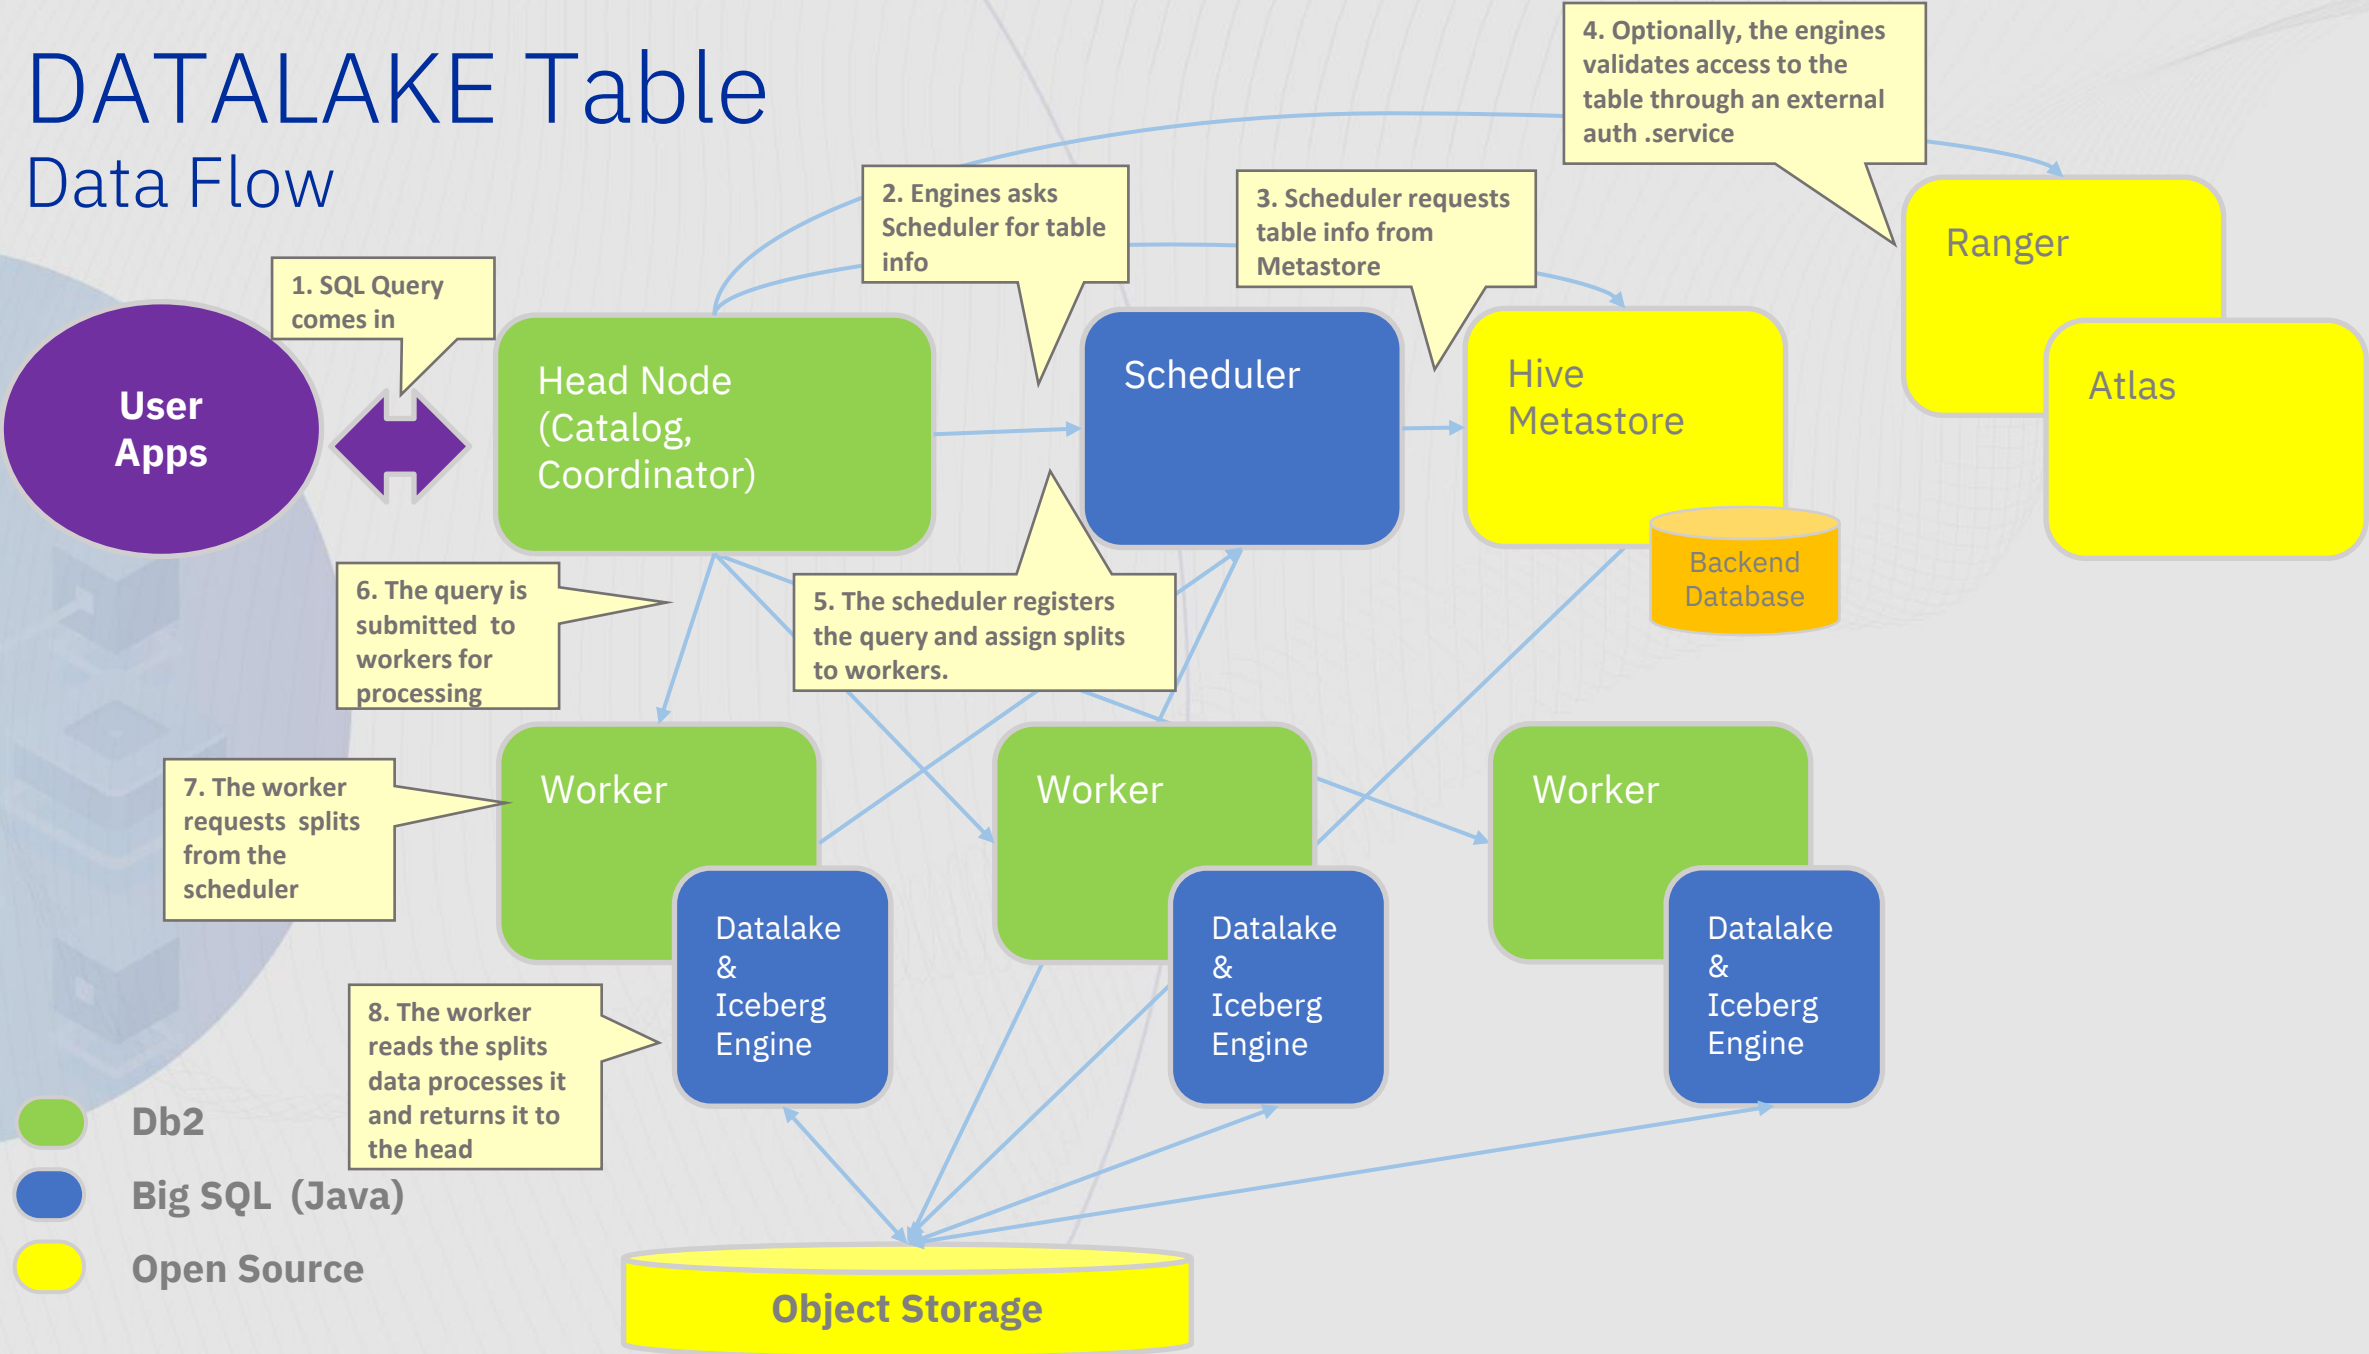
IBM Db2 Warehouse

DATALAKE Table Support



DATALAKE Table

Data Flow



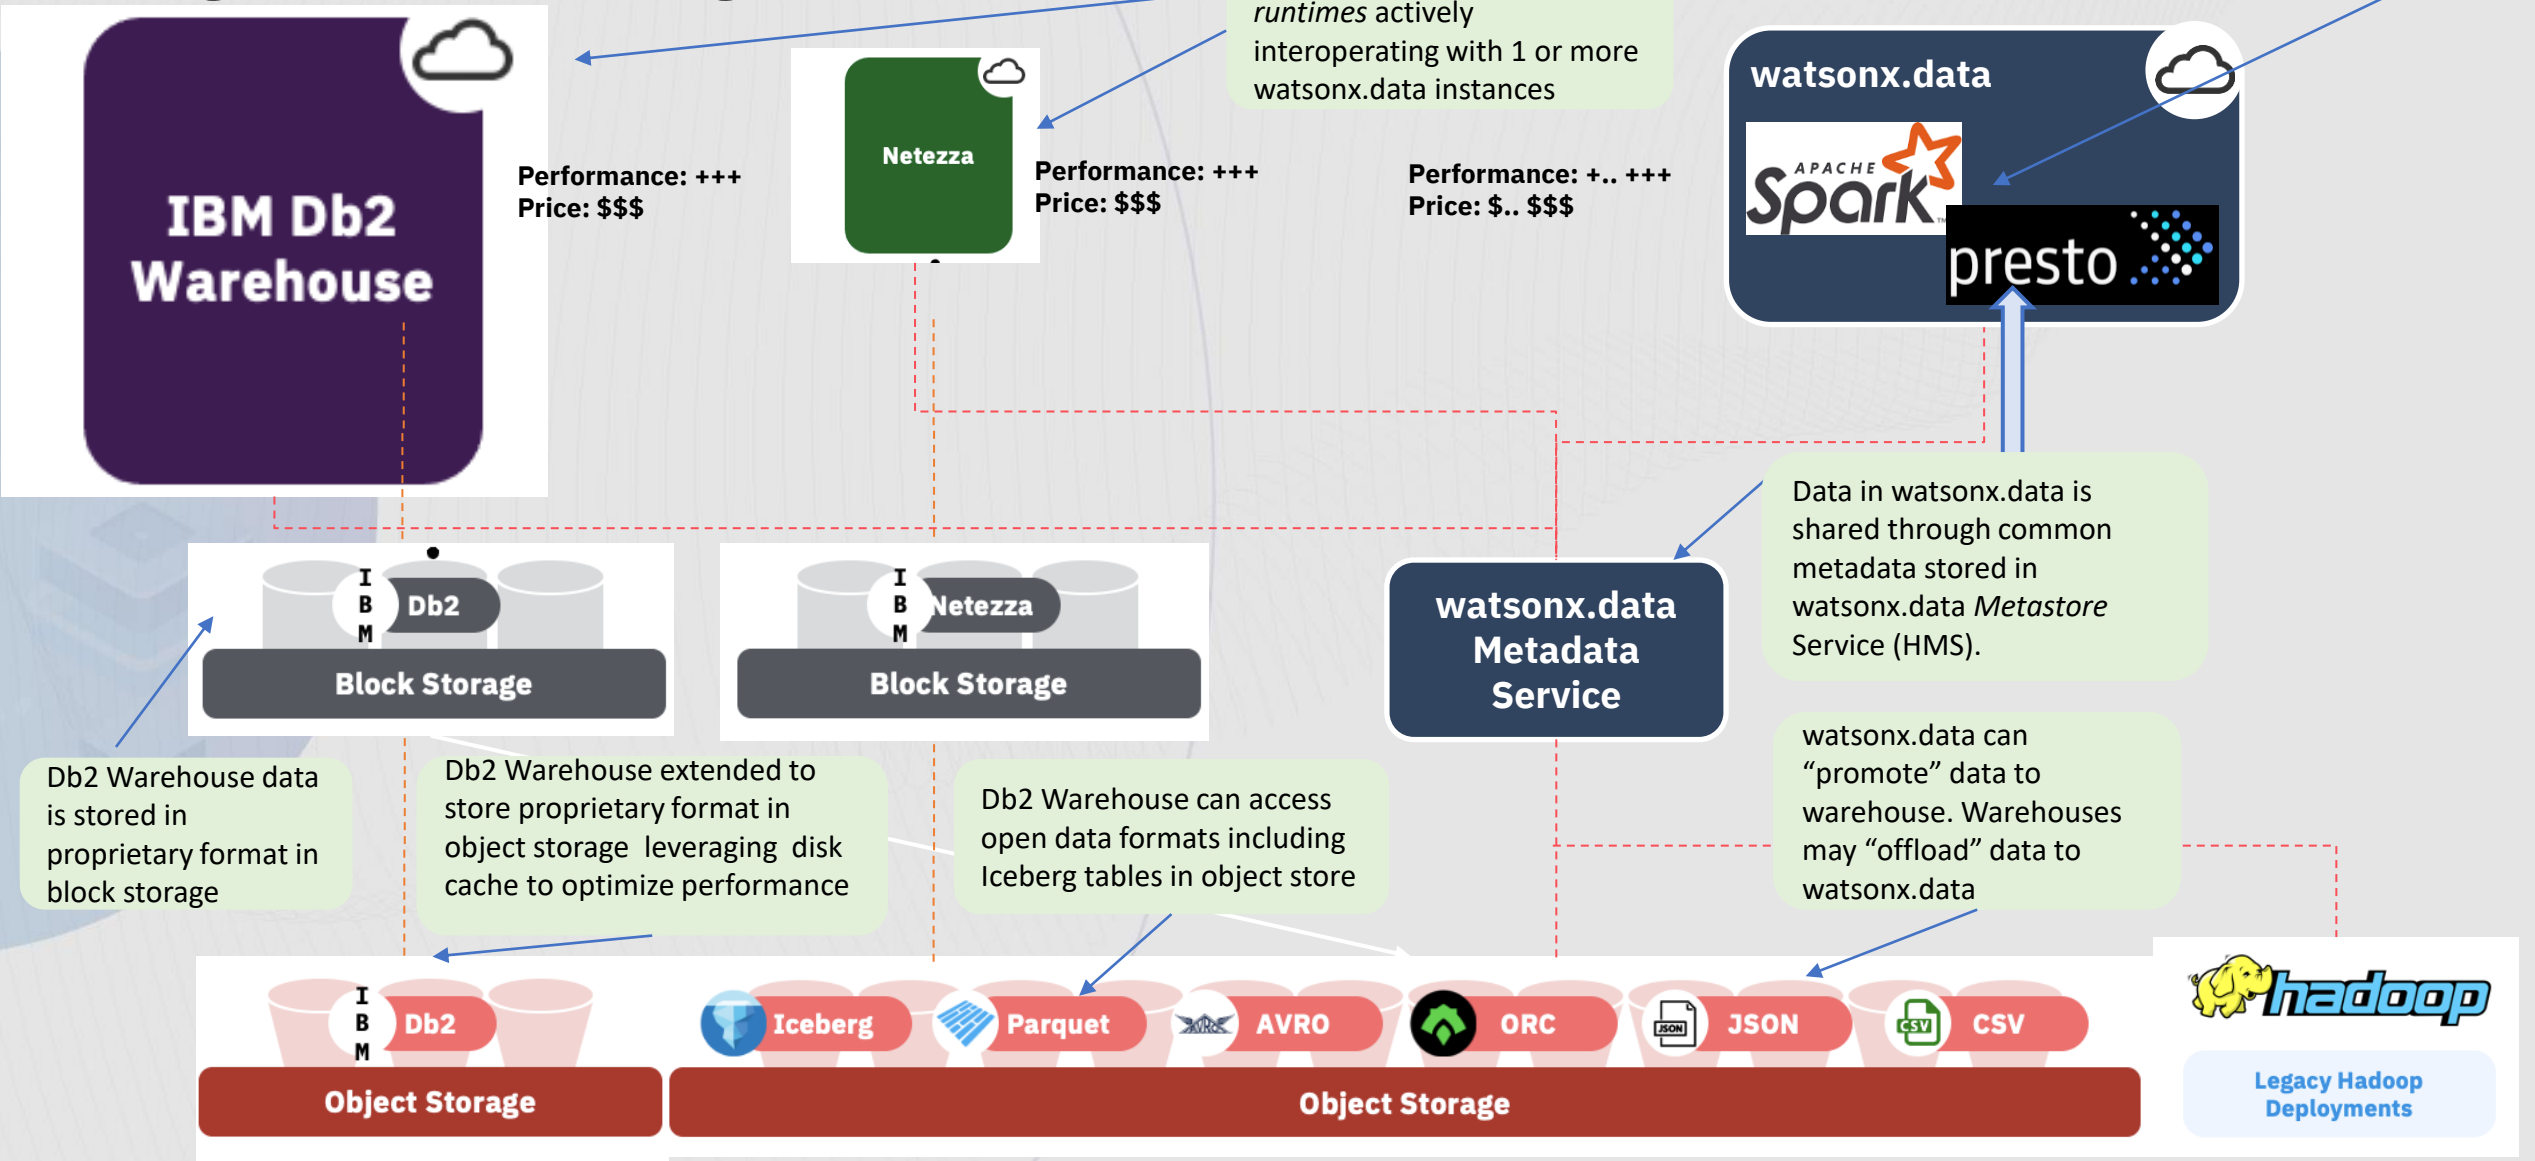
AGENDA

Backup Deep Dives
- watsonx.data integration



Better together watsonx.data & Db2 Warehouse

Designed to work together



AGENDA

2023 Deliverables

11.5.9 Overview

Deep Dives

Things starting with CS

Backup Deep Dives (skipped as covered in detail next 2 days)

Reference: 11.5.0 to 11.5.8 Overviews

Db2 11.5.8 Highlights



Hybrid & Multi-cloud

Containerization:

- RHOS on AWS & Azure - ROSA and ARO
- Deploy on Amazon EKS
- Deploy on Azure AKS
- Integration with Velero for BAR

- **Note: most of above not tied to 11.5.8 but on similar 2H 2022 timeline (see roadmap for details)**

- Operator driven Db2 native backup and restore
- Db2 W Scaling - vertical by adjusting resources per pod
- Db2 W Scaling - horizontal by scaling out number of pods
- HADR role-aware Kubernetes service for Db2U. Route clients to primary.



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
 - db2dsdriver validation tool
 - Large result set query perf
 - .NET Visual Studio 2022 support
 - OpenJDK 17 runtime support

* Coming soon (1Q 2023)



Columnar Enhancements

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



Cognitive and Emerging

Db2 Graph, new tutorials plus:

- GUI Enhancements to Query Screen
- Discovery mode to explore w/o writing gremlin queries
- History to revert and compare graphs models
- Autogenerate code to run queries

Db2 11.5.8 Highlights



Availability

- **db2adutl 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
- GUI 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

- pureScale TCPIP on AWS – deployment and support (already available).
- **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
- Avoid lock timeout due to EHL
- pureScale on z15 support (not be tied to modpack, will be later)



Core Engine

- 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.
- New ADMIN_FMP_TERM SP to shutdown an fmp process
- External table support to read from “binary numeric” format
- Extend db2RemStgManager to use storage aliases
- **Support for TLS 1.3**
- **Force Load of GSKIT which comes with db2**
- **Monitor improvements: SSL info to be added to MON_GET_CONNECTION**
- Catalog And Process Model Serviceability improvements
- db2support to automatically gather text search logs
- GUI and silent installs now can install Pacemaker
- Java, TSA, GPFS and Pacemaker updates
- Windows 2022

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

- AIX 7.3 / 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

Db2 11.5.6 Highlights



Cognitive and Emerging

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



Hybrid & Multi-cloud

RHOS Container Updates

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



Application Development

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



Install / Deployment

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



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

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 Import/Export
- Updates to default port numbers to avoid Linux Ephemeral port
- LOCK NO WAIT and WAIT time on Select/update/delete

installFixpack now license acceptance 11.5.6 and on



Heads
up!

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

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
 - PostgreSQLV10 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 11.5.4 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 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