Db2 11.5.9 – Next Generation Data Warehousing

Les King
lking@ca.ibm.com
July 2024
Db2 Night Show

Agenda

- Strategic Drivers
- watsonx.data
- Remote Tablespace Support (Native Cloud Object Store)
- DATALAKE Table
- Db2 Warehouse SaaS Gen 3
- IIAS Follow-on
- A Peak Ahead
- Survey Questions

Strategic Drivers

Key: EKS AKS nShift astructure

IBM **Db2**

Portfolio of database solutions

Built to run the world's mission critical workloads

Cloud/SaaS

Db2

SaaS

—

Relational database delivered as a service

IBM Cloud

Cloud/SaaS

Db2 Warehouse

SaaS

—

Cloud data warehouse delivered as a service

IBM Cloud aws

Software

Db2

—

Relational database built to run the world's mission critical workloads

Software

Db2 Warehouse

—

High-performance data warehouse for deep analytics and machine learning

Software

Db2 BigSQL

—

SQL-on-Hadoop engine, delivering MPP and advanced data query

Evolving the Needs of All Data Professionals

As data maturity increases, so does the number of data professionals who are hungry to put data to work

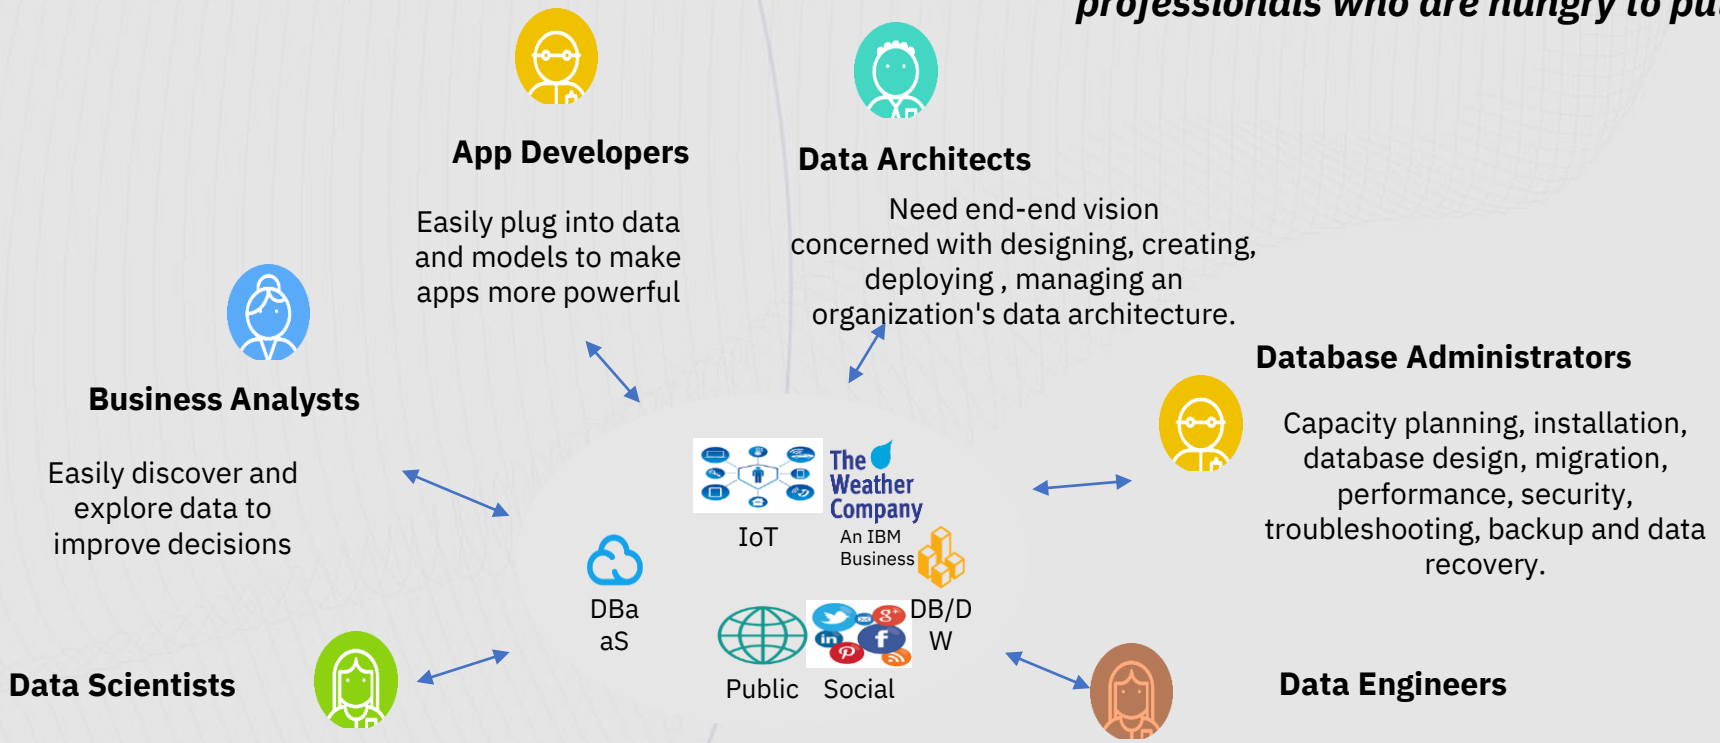

Db2 V11.5 – Balance between foundation and modernization

Rock Solid Database

Infuse AI

Containers

Making Data Simple

Enterprise Readiness

- *Performant*
- *Secure*
- *Available*
- *Automated administration & monitoring*

Cost Savings

- *ML optimizer*
- *Adaptive workload management*

Making Data Accessible

Modern Development

- *Developer-friendly*
- *Multi-model*

Consumability

- *Deploy across multiple form factors*
- *Multi Cloud and Containers*

Db2 and AI

A prescriptive approach to accelerating the journey to AI

AI

INFUSE – Operationalize AI with trust and transparency

ANALYZE – Scale insights with AI everywhere

ORGANIZE – Create a trusted analytics foundation

COLLECT – Make data simple and accessible

Data of every type, regardless of where it lives

MODERNIZE your data estate for an AI and multicloud world

Db2 – Ready for Modern Deployments and Workloads

Containerized Deployment Options

The ability to deploy Db2 leveraging containerized platforms including RHOS and non-RHOS platforms

Fully Managed DBaaS

The ability to deploy Db2 as a fully managed service for public, multi and hybrid cloud environments

Cornerstone of COLLECT in Journey to AI

Db2 is the cornerstone for COLLECT in Cloud Pak for Data to cover end-to-end needs to leverage AI

Self Managed Db2 Deployments

The ability to deploy proven reference architectures of Db2 – for any workload – on all key cloud providers

Data Virtualization

Db2 contains a data virtualization component which allows Db2 to be a doorway to all of your business data

In-Db2 Machine Learning

Allows data scientists and developers to bring machine learning local to the data stored within Db2

Multi-Model – NoSQL and NewSQL Data Store

Db2 is a multi-model data store supporting native relational, JSON, BSON, Graph, Spatial, Text and XML

Mixed Workloads

Db2 can handle any combination of workloads including real-time data ingestion, multi-model and mixed.

Db2 11.5.9

Db2 V11.5 - Lifecycle

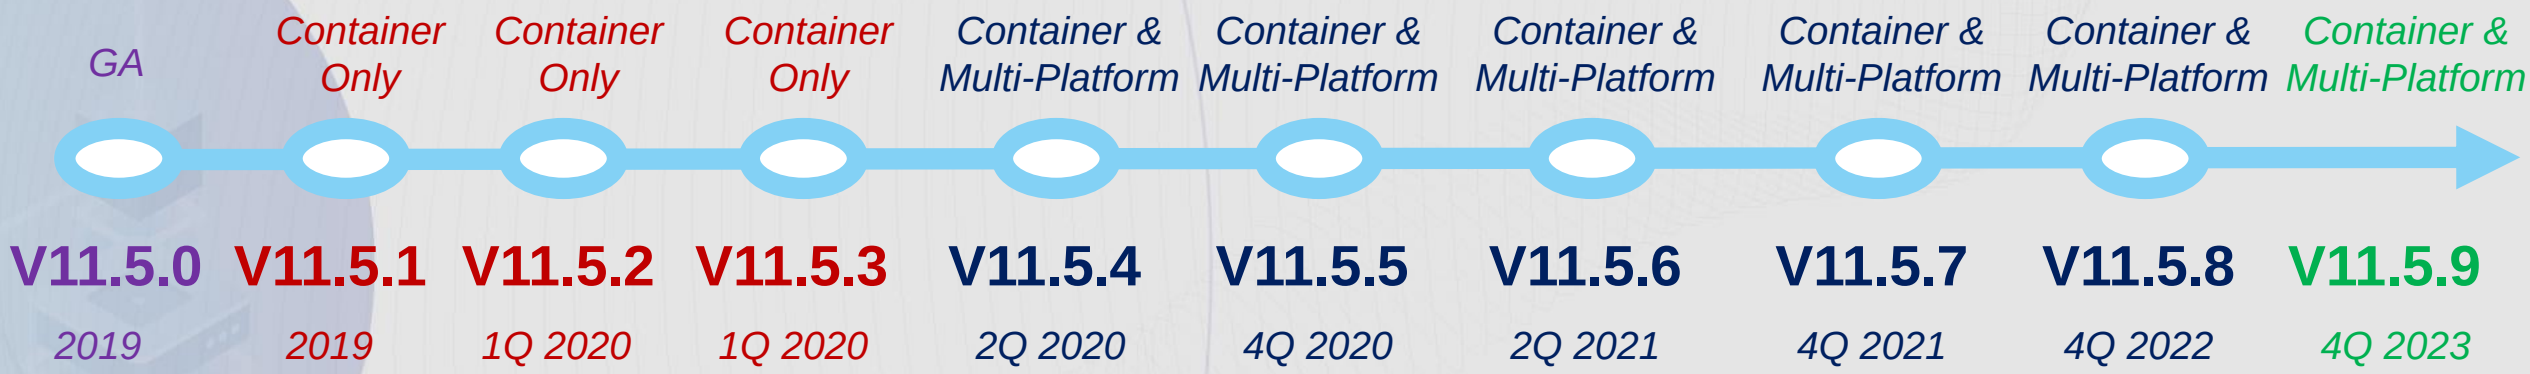

Regular fix/mod pack deliveries

- With overriding focus on stability for immediate production deployment
- Modifications will contain select functional enhancements (often off by default)

Db2 Calendar

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)

IBM watsonx

The platform for AI and data

watsonx

Scale and accelerate the impact of AI with trusted data.

watsonx.ai

Train, validate, tune and deploy AI models

- A next generation enterprise studio for AI builders to train, validate, tune, and deploy both traditional machine learning and new generative AI capabilities powered by foundation models. It enables you to build AI applications in a fraction of the time with a fraction of the data.

watsonx.data

Scale AI workloads, for all your data, anywhere

- Fit-for-purpose data store optimized for governed data and AI workloads, supported by querying, governance and open data formats to access and share data.

watsonx.governance

Enable responsible, transparent and explainable data and AI workflows

- End-to-end toolkit encompassing both data and AI governance to enable responsible, transparent, and explainable AI workflows.

The platform for AI and data

watsonx

Scale and accelerate the impact of AI with trusted data.

watsonx.ai

Train, validate, tune and deploy AI models

- A next generation enterprise studio for AI builders to train, validate, tune, and deploy both traditional machine learning and new generative AI capabilities powered by foundation models. It enables you to build AI applications in a fraction of the time with a fraction of the data.

watsonx.data

Scale AI workloads, for all your data, anywhere

- Fit-for-purpose data store optimized for governed data and AI workloads, supported by querying, governance and open data formats to access and share data.

watsonx.governance

Enable responsible, transparent and explainable data and AI workflows

- End-to-end toolkit encompassing both data and AI governance to enable responsible, transparent, and explainable AI workflows.

Market Dynamics

Major disruptions are driving the growth in the analytics repositories market **from on-prem to SaaS** and **from proprietary to open technologies**

Analytics Repositories Market Landscape

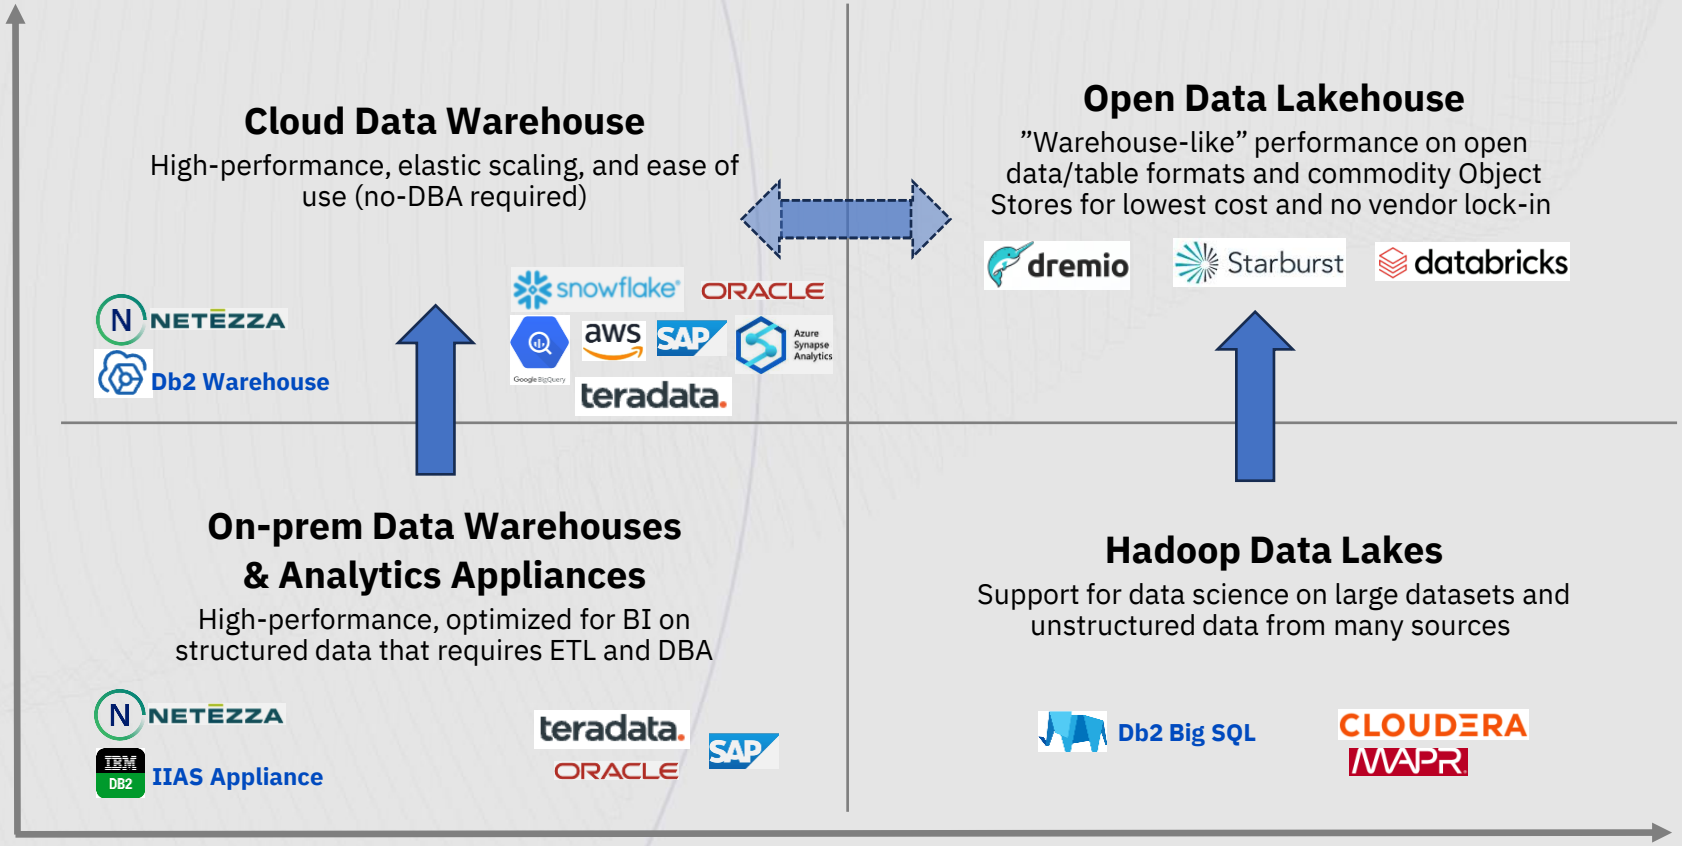

SaaS
$31bn 2025
27% CAGR ('21-'25)

Deployment

On-prem
$12bn 2025
2% CAGR ('21-'25)

Proprietary
$26bn 2025
13% CAGR ('21-'25)

Technology

Open
$17bn 2025
27% CAGR ('21-'25)

Sources: IDC Data Management Forecast (November 2021), IDC BDA Forecast (June 2021), MI modeling

The Data Lakehouse

The Data Lakehouse implements the **data structures and management features** of a data warehouses on the **low cost, reliable & scalable** object storage within a new architectural approach that leverages open-source technology.

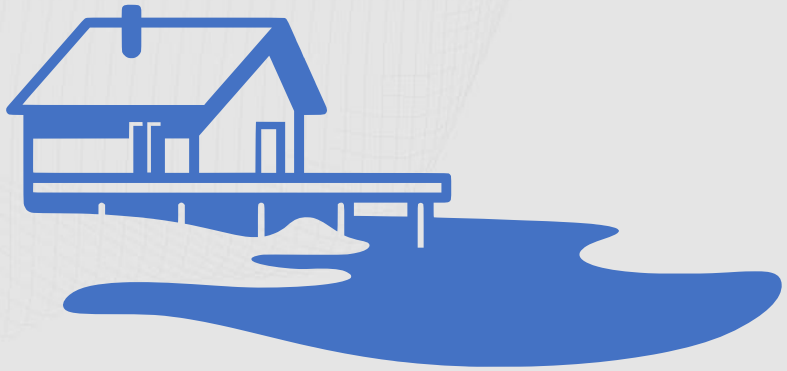

It enables organizations to manage their data in an **open, flexible, cost-effective, feature rich and scalable way**, enabling Business Intelligence and Machine Learning on all data.

Data *Lake*
+
Ware*house*

The **IBM** Data Lakehouse

The IBM data lakehouse brings together the advantages of data warehouses and data lakes within a new managed cloud service and self-managed on any Cloud or on-premise.

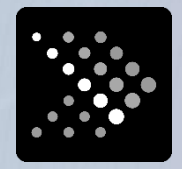

A Low-Cost and extensible Query Engine
Presto is an **open-source, fast** and reliable **SQL engine** for Data Analytics and data lake houses.

A proven and reliable metadata repository
The hive **metastore** is the de facto standard in open-source data lake metadata management

Stores data in Object Store buckets in the Iceberg open data format to **facilitate data access and sharing** across applications

Open, Flexible, and Modular
Lakehouse is designed to enable customers to standardize their data formats and metadata with unified data governance

Evolve your big data platform
The simplest path to upgrade from traditional big data platforms, either as a side car or moving data to cloud object storage

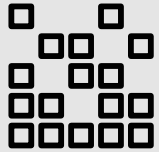

Limitless scalability and elasticity
Explore, shape, and analyze data at any scale by separating storage and compute

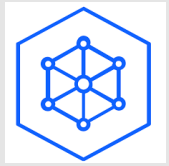

Integrates readily with Db2 Warehouse and Netezza to support the **right data engine for the right workloads at the right cost**

Better together Lakehouse & Warehouse

Designed to work together

Any number of these *runtimes* actively running inside Lakehouse instance

Any number of these *runtimes* actively interoperating with 1 or more Lakehouse instances

IBM Db2 Warehouse

Performance: +++
Price: $$$

Netezza

Performance: +++
Price: $$$

Performance: +.. +++
Price: $.. $$$

IBM Lakehouse

Apache Spark

presto

Warehouse data is stored in proprietary format in block storage

IBM Db2 Block Storage

IBM Netezza Block Storage

Lakehouse Metadata Service

Data in the Lakehouse is shared through common metadata stored in Lakehouse *Metastore* Service (HMS).

Warehouse extended to store proprietary format in object storage leveraging disk cache to optimize performance

Warehouse can access open data formats including Iceberg tables in object store

Lakehouse can "promote" data to warehouse. Warehouses may "offload" data to the lakehouse.

IBM Db2 Object Storage

Iceberg Parquet AVRO ORC JSON CSV Object Storage

hadoop

Legacy Hadoop Deployments

DATALAKE TABLE Support

Datalake Tables

- A Data Lake “Table” is a collection of files serialized using an **Open Data File** (ODF) format (CSV, ORC, Avro, Parquet …) stored on remote storage (HDFS, S3, COS, …)

- The **metadata** of the table is stored in a Metastore server
  - Location
  - Schema
  - Partition lists

- An engine querying the table must query the metadata first and can proceed to read the data from remote storage

- Benefits
  - Interoperability of open data formats
  - Ease of use

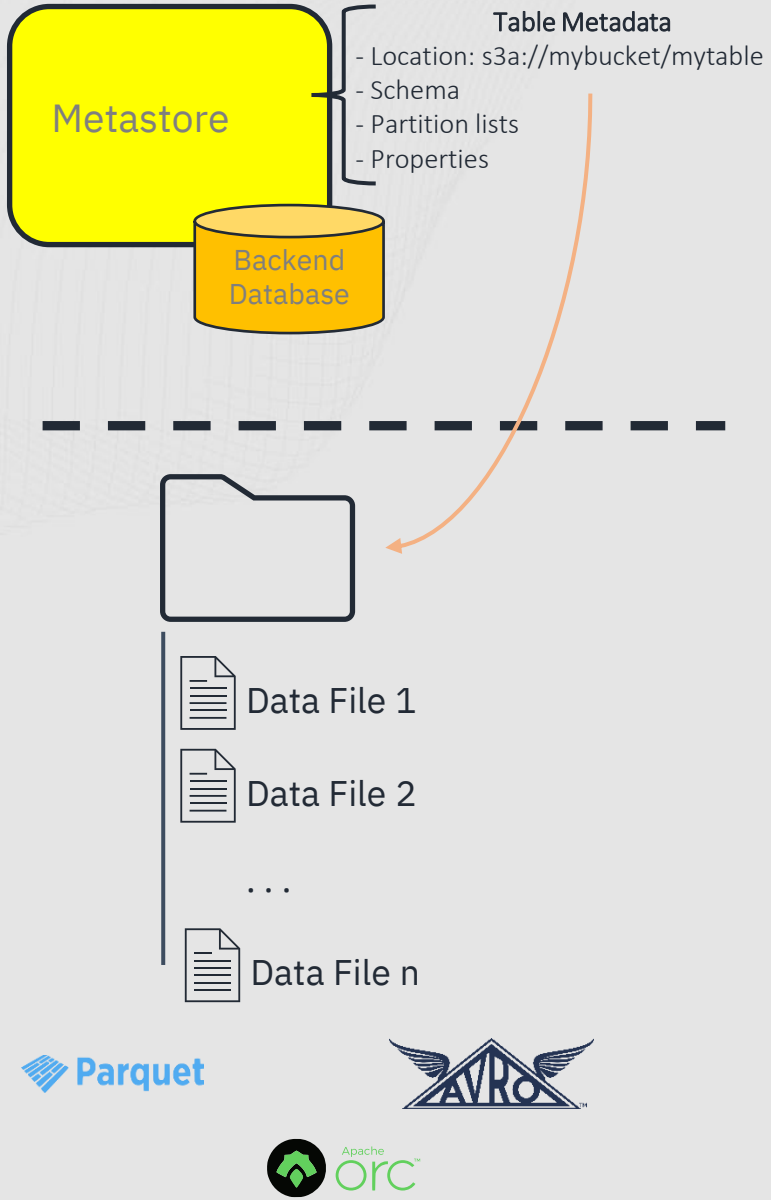

IBM Db2 Warehouse

Datalake Table Support

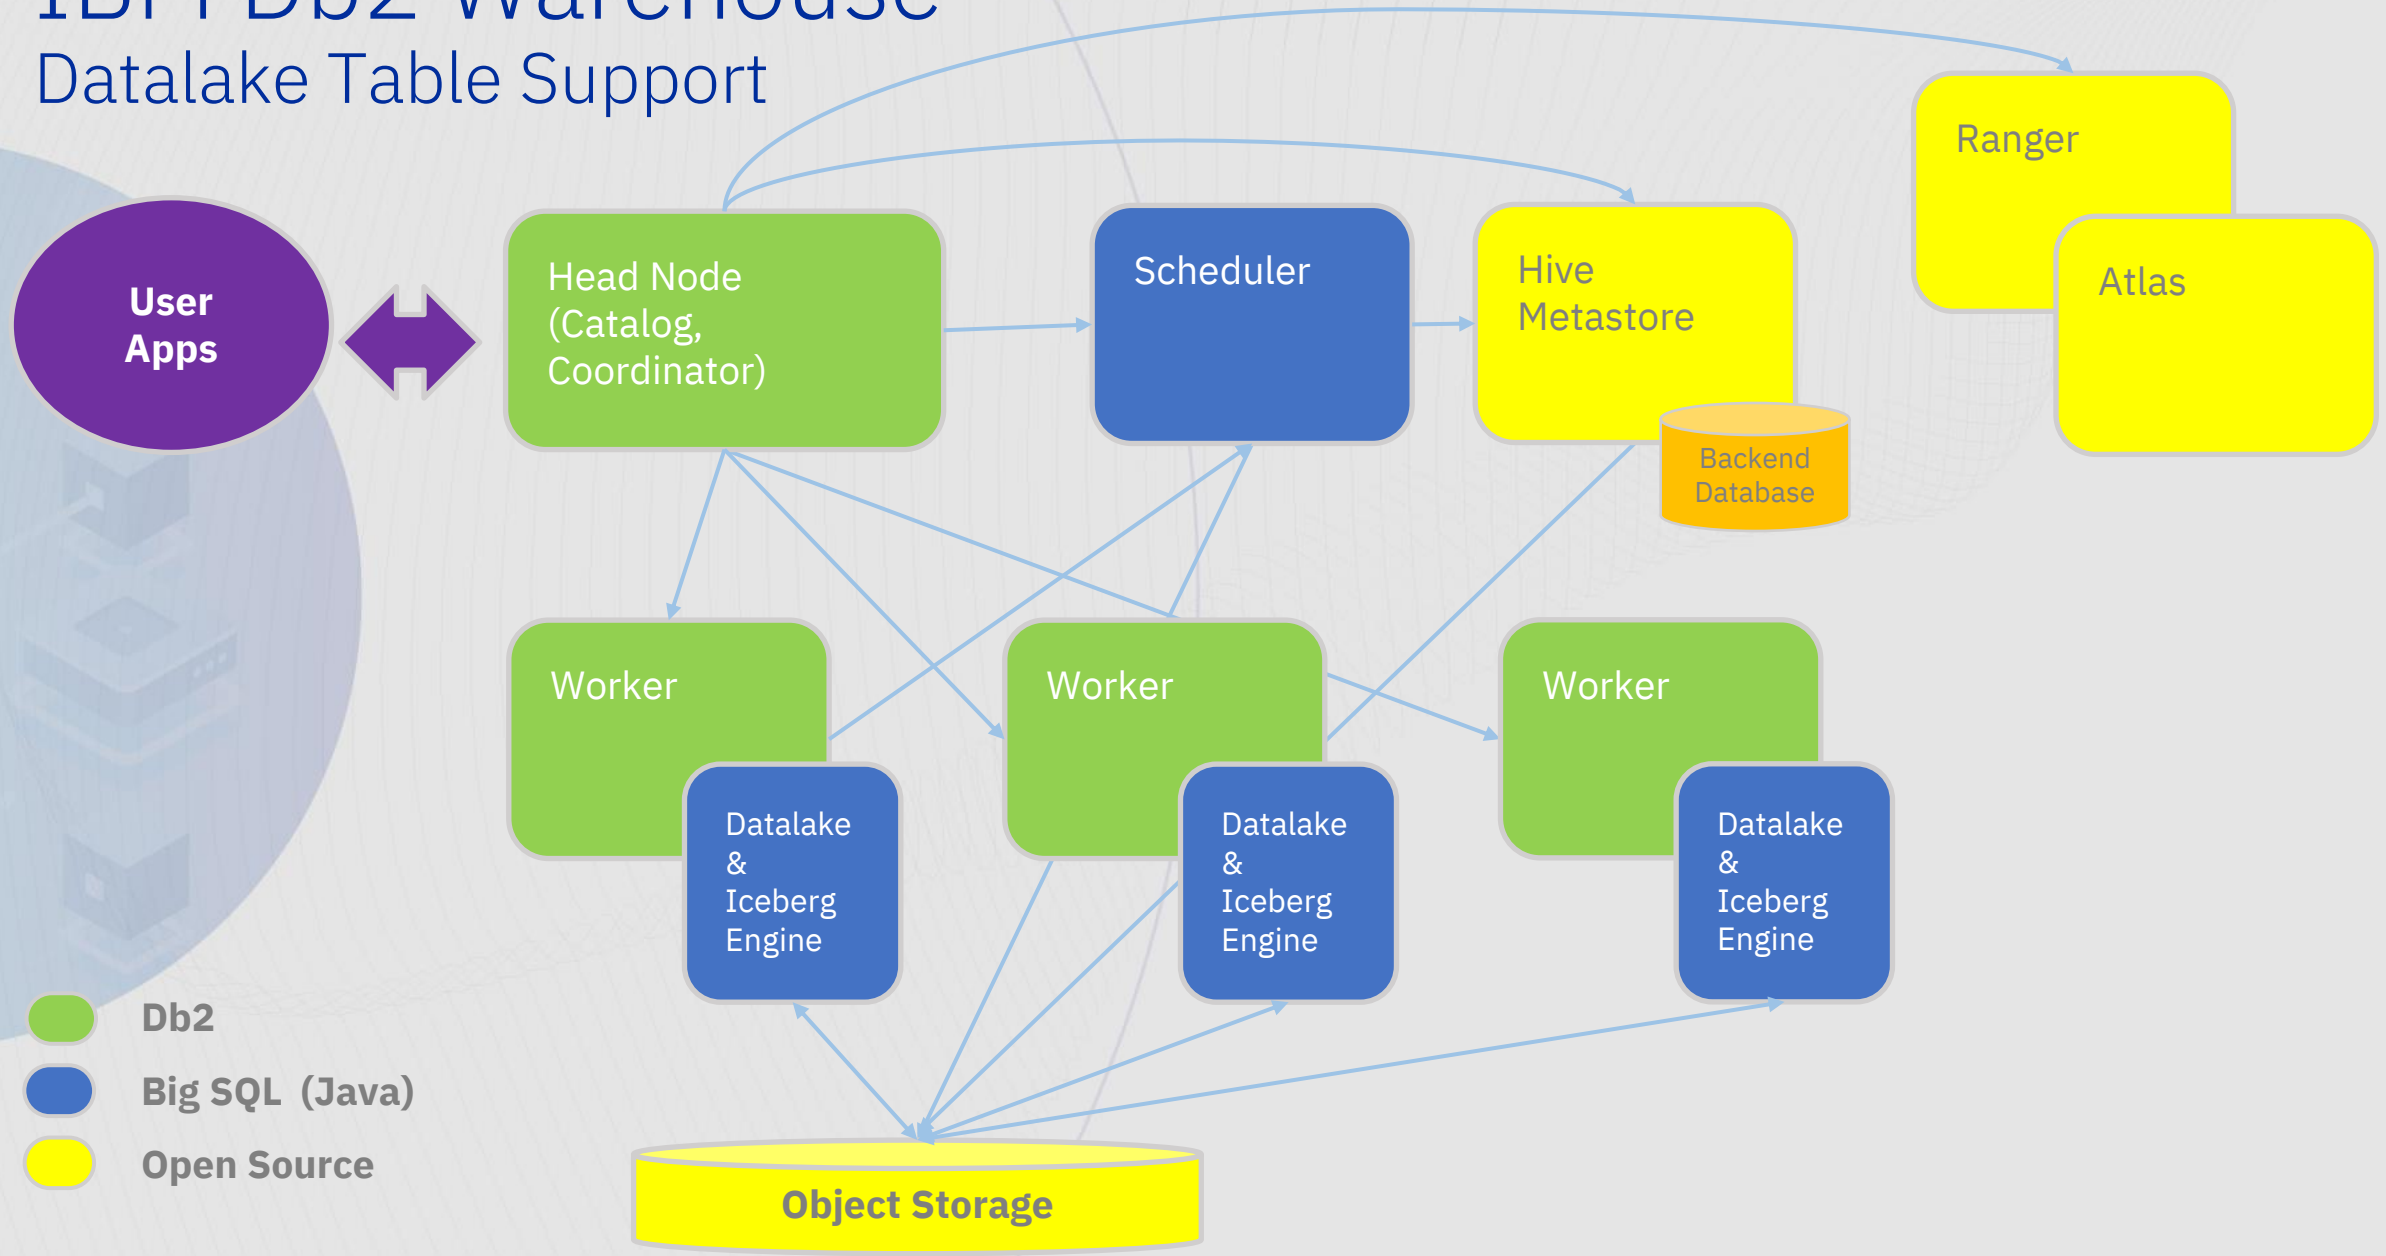

Open Data File Format Limitations

- A Data Lake “Table” is a collection of files serialized following an Open Data File format

- Passive data structures – serialized set of data records
  - No notions of their **state** or **history**
  - No concurrency control between applications
  - **No ACID**, even less transactions

- Separate metadata
  - Need for a “**Catalog**”
  - No awareness of catalogs – it’s an external system

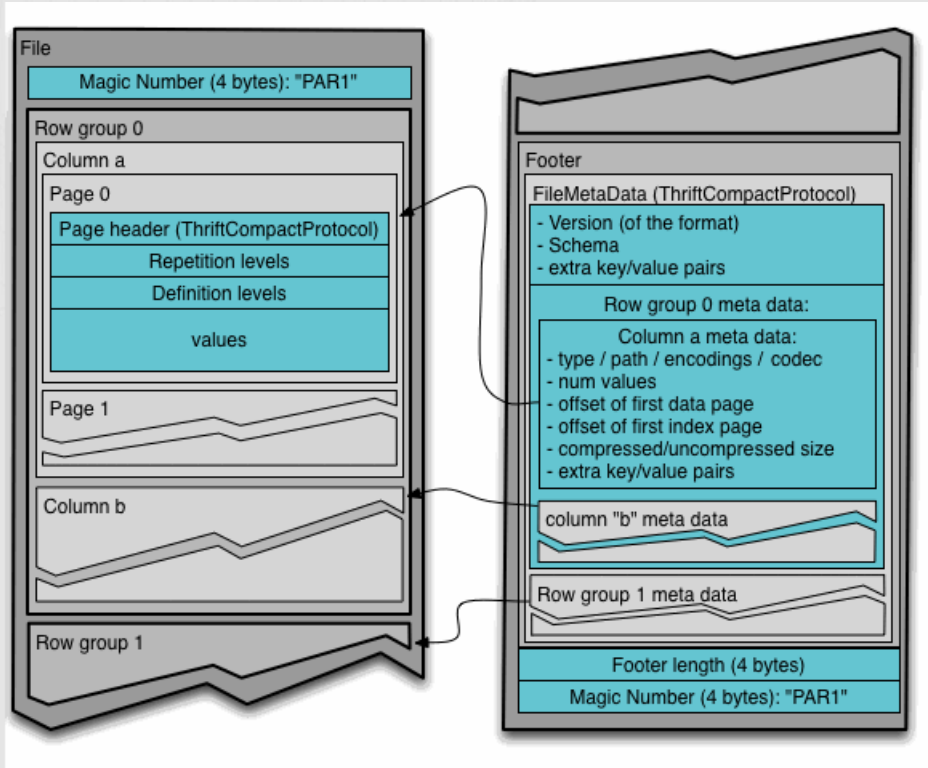

Parquet

Apache orc™

AVRO™

Apache Iceberg

An Open Data Table format for the Lakehouse

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

Support for CREATE, SELECT & INSERT including partitioning support

No UPDATE, DELETE

No Scheme Evolution

No Time Travel

Smaller restrictions related to Icebert/Db2 type compatibility such as nested types, etc.

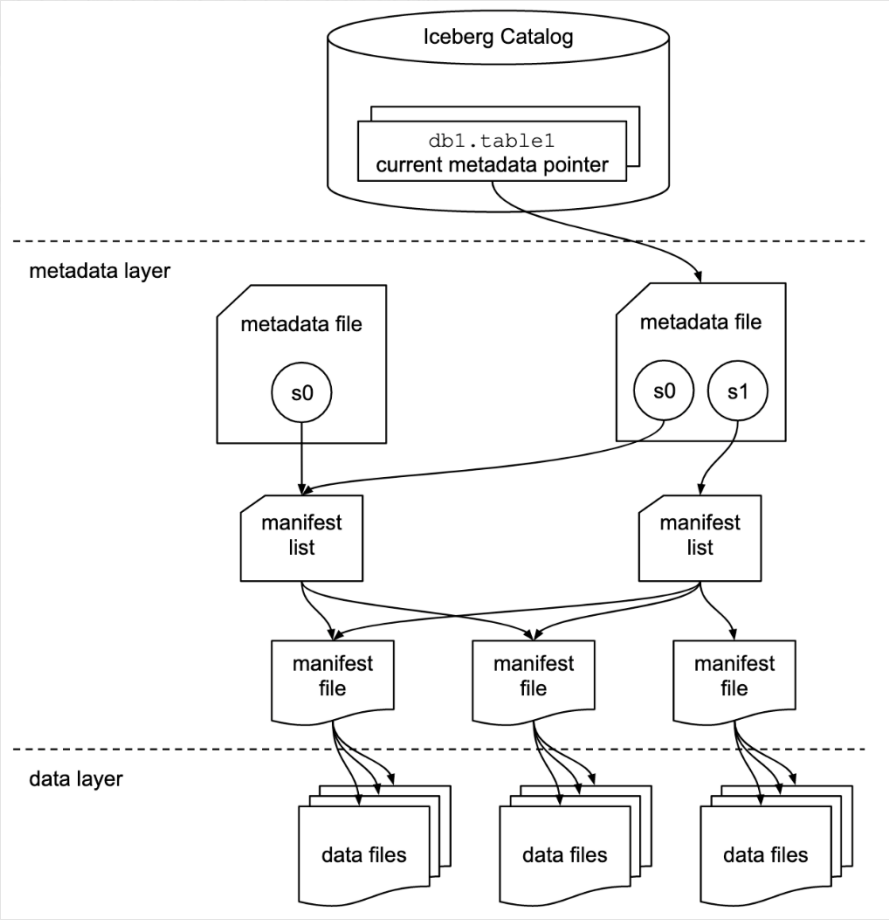

REMOTE TABLESPACE Support

Object Storage support for Table Storage – Remote Tablespace Support

- Db2 now supports Amazon S3 object storage for database table storage, where customer data resides within the database
- Customer saves cost by using object storage instead of block storage
  - Allows customers to store data across a mix of block storage and object storage, based on business or technical requirements
  - Db2 uses different mechanisms in order to facilitate reads and writes to object storage
- No applications and workload changes necessary in order to use this feature
  - Db2 handles all the necessary interfacing to object storage, thus existing applications and warehouse workloads do not have to be changed in order to make use of this
  - Specific Db2 tablespace available backed by S3 for customer use
  - Insert, Update, Delete data as needed into and out of tables within object storage
  - Move and copy data to and from column-organized tables residing in block storage and object storage
  - Query data seamlessly no matter where it resides (in block or object storage), in isolation or in combination with each other

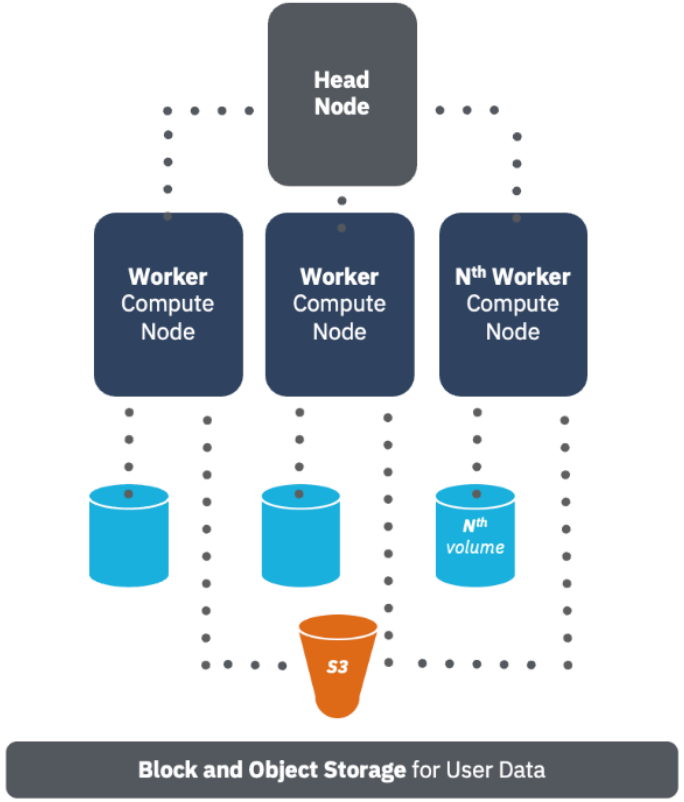

Up to 34x lower storage cost when compared against the prior generation

34x

Less expensive to host Db2 data on object vs block storage[1]

Costs of Db2 Warehouse block vs object storage.

Db2 Warehouse on Cloud (SaaS) – Gen 3

Db2 Warehouse (SaaS) – Gen 3 – Based on Db2 11.5.9

Fully managed / SaaS

Focus on the analytics, we'll take care of the rest

Blazing-fast

Columnar-organized, memory-optimized data warehouse

Scalable & elastic

Independently scale and manage compute & storage

Continuously available

Managed compute, highly available storage, cross-cloud replication

Reliable

Double protection with disaster recovery & self-service backup/restore

Secure

Encrypted at the storage level by the vendor and customer can protect data using their Key Protect keys.

Deploy as fully-managed DBaaS on **IBM Cloud** and **Amazon Web Services**

Db2 Warehouse (SaaS) – Gen 3 – Cloud Object Store (Db2 11.5.9)

- Db2WoC now supports Amazon S3 object storage for database table storage, where customer data resides within the database
- Customer saves cost by using object storage instead of block storage
  - Allows customers to store data across a mix of block storage and object storage, based on business or technical requirements
  - Db2WoC uses different mechanisms in order to facilitate reads and writes to object storage
- No applications and workload changes necessary in order to use this feature
  - Db2WoC handles all the necessary interfacing to object storage, thus existing applications and warehouse workloads do not have to be changed in order to make use of this
  - Specific Db2 tablespace available backed by S3 for customer use
  - Insert, Update, Delete data as needed into and out of tables within object storage
  - Move and copy data to and from column-organized tables residing in block storage and object storage
  - Query data seamlessly no matter where it resides (in block or object storage), in isolation or in combination with each other

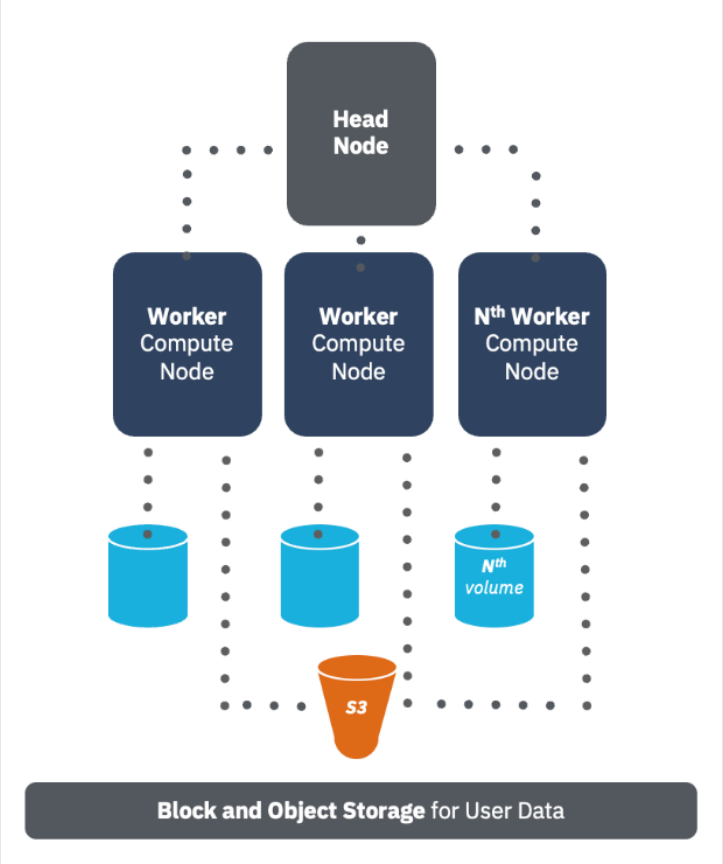

Db2 Warehouse (SaaS) – Gen 3 – DATALAKE Table (Db2 11.5.9)

- Db2WoC now also supports Open Data Formats as DATALAKE tables, allowing for seamless access to other data within the enterprise for integrated workloads.

- Leverage existing compute resources dedicated to the warehouse
  - Facilitate data use to and from the Db2 Warehouse to quickly access a variety of enterprise data
  - Leverage the high performance Db2 engine for queries against enterprise data

- Db2WoC provides interfaces for customers to leverage their enterprise data residing in object storage as DATALAKE tables
  - Supports both regular and Iceberg DATALAKE table types, based on existing data formats or for business/technical requirements such as ACID compliance
  - Browse, explore, and query enterprise data in both Db2 and DATALAKE formats, using either the web-based UI, or through SQL
  - Access data in place within DATALAKE tables, joining as necessary with Db2 based data for queries
  - Access data within DATALAKE tables and import into Db2 formatted tables
  - Create new DATALAKE tables in S3 and export from Db2 formatted tables

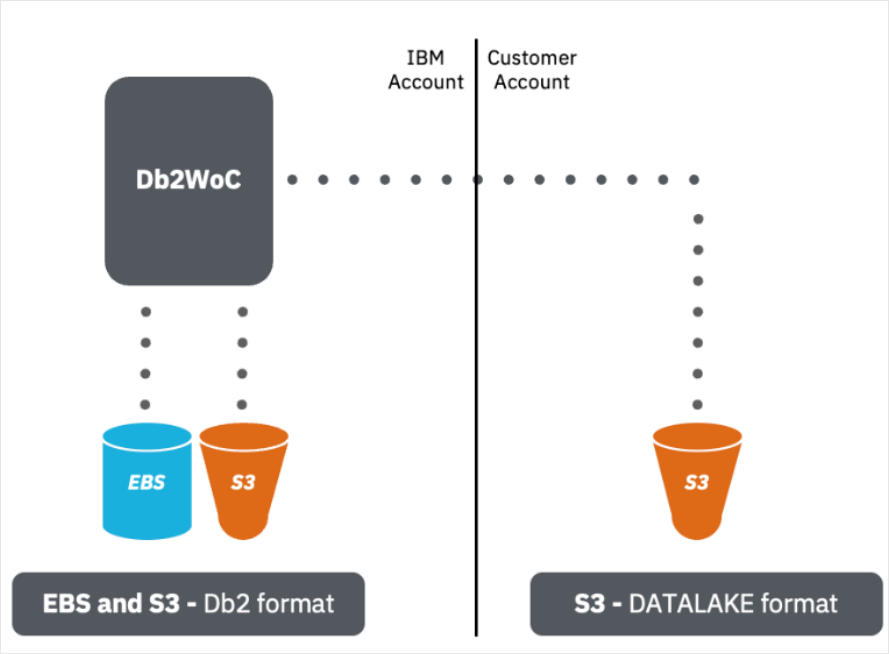

Db2 Warehouse (SaaS) – Gen 3 – Based on Db2U

- Managed via k8s Custom Resource (CR)
  - Kind: `Db2uInstance`
  - Short name: `db2`
- Underlying Kubernetes resource model:
  - Db2 Engine Pod lifecycle managed using a NEW `Db2uEngine` CR, which manages a collection of Db2 Pods.
  - Onetime tasks managed via a `Job` resource
  - In-pod HA to recover Db2 failures, avoiding a pod lifecycle event. This built-in HA leverages ETCD for state information
  - Lifecycle of (stateless) Add-Ons (REST, Graph, Q-rep, etc.) managed via `Deployment` resources

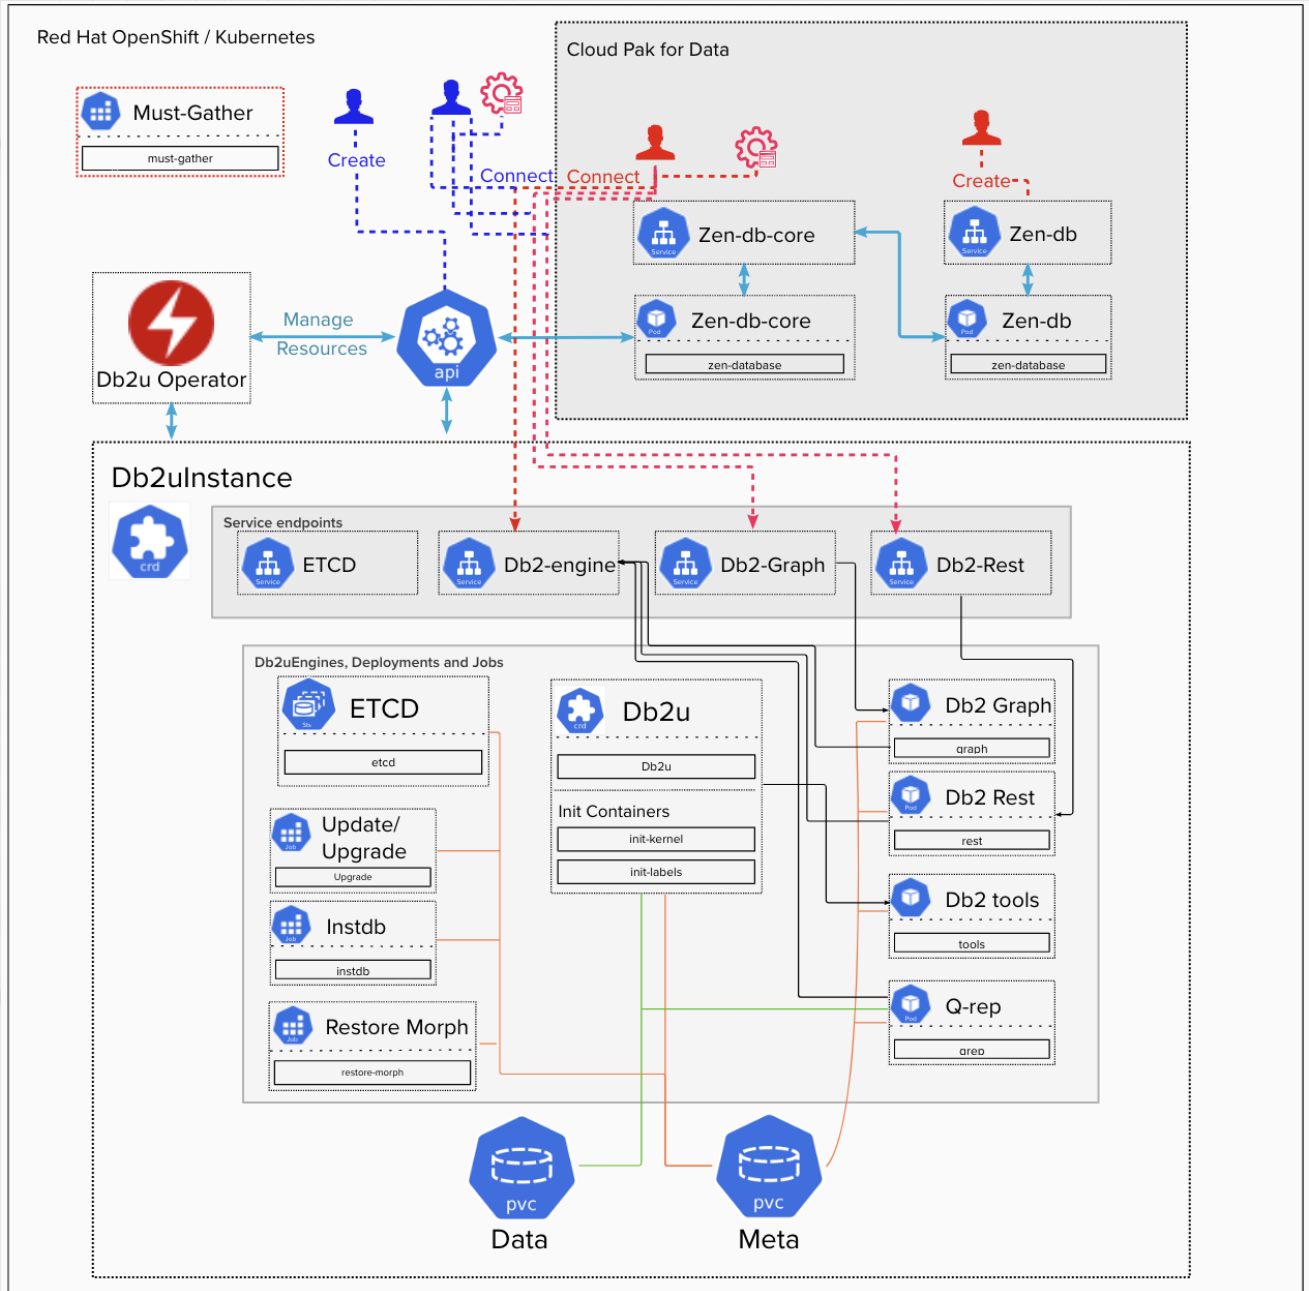

Db2 Warehouse (SaaS) – Gen 3 – Highly Scalable

- Scale using the Console or APIs
  - The platform provides mechanics to initiate a scale (we create a scale recipe)
  - In the scale recipe, we provide the new core and storage values.

- The operator framework will
  - Gracefully shutdown the engine and scale down the formation
  - Bring it up in the new spec (appropriately setting the correct MLN count per pod)
  - As the new spec is brought up, new nodes are automatically added by the Cluster Auto Scaler
  - If it is a shrink operation, unused nodes are automatically returned by the Cluster Auto Scaler

- Compute Expansion / Shrink incurs a Downtime

- Storage scaling works the exact same way (recipe mechanism)
  - Storage scale up is a totally online operation
  - Block Storage shrink (support date TBD) will be offline in nature

Db2 Warehouse (SaaS) – Gen 3 – Improved HA

- **Multi prong, provided by**
  - Wolverine (our HA component) - Core Db2 container
  - Kubernetes - non Db2 containers
- **Wolverine responsible for and provides**
  - System recovery (when the pods fail)
  - Overall System health and status recording in ETCD
  - Device Manager framework
  - HA loop (for monitoring and handling HA for processes within core Db2 Container)
  - API endpoints
- **Kubernetes**
  - Pod level HA based on liveliness probes

- **When a Db2 node goes down,**
  - Kubernetes will automatically schedule the pod to run on a free node from the pool
  - Wolverine will wait for the pod to come back up and take the engine through a recovery process
  - Engine is back running in full capacity within a short time.
- **When a Db2 pod goes down**
  - Kubernetes will restart the pod
  - Wolverine will wait for the pod to come back and take the engine through a recovery process.

Db2 Warehouse (SaaS) – Gen 3 – More Granular Backup/Restore

- Backups - a combination of Snapshot Backups and S3 backup
  - Scheduled every 24 hour period or on demand
  - Scheduled or on-demand both internally create a Recipe
  - The Recipe wil go through the following phases
    - Copy s3 metadata to block storage location
    - Do a db2 write suspend (including S3 deletes off)
    - Create volume snapshots
    - Initiate S3 backup process
    - Backup metadata
    - Write resume (but leave S3 delete on)
    - When AWS S3 backup completes (tag backups and turn on S3 deletes)

Restore process, On-demand, internally creates a Recipe.
The recipe will go through the following phases

- Stop Db2
- Initiate meta data restore
- Initiate S3 restore
- Initiate volume snapshot restore
- When S3 restore is ready to use, scales all pods up

Moving forward from IIAS

Moving Forward from PDOA/IIAS

IBM Integrated Analytics System (IIAS)

IBM PureData System for Operational Analytics (PDOA)

Migrate to...

1. Db2 Warehouse SaaS
Fully managed, high performance, cloud-elastic data warehouse
↳ Runs on IBM Cloud, AWS
↳ Available in US, EU, UK, JP, AP commercial regions

2. Db2 Warehouse Self-managed
Self-managed, data warehouse on cloud provider managed Kubernetes or OpenShift infrastructure.
↳ Runs on Amazon EKS, Amazon ROSA, Azure AKS, Azure ARO
↳ Available worldwide

3. Db2 Warehouse Cloud Rack
On-premises, OpenShift high performance, data warehouse deployable on Power 10 or X86 infrastructure
↳ Runs on POWER10 or x86 reference architecture

Moving Forward from PDOA/IIAS – Cloud – Fully Managed

IBM Integrated Analytics System (IIAS)

IBM PureData System for Operational Analytics (PDOA)

Migrate to...

1. Db2 Warehouse SaaS
Fully managed, high performance, cloud-elastic data warehouse
↳ Runs on IBM Cloud, AWS
↳ Available in US, EU, UK, JP, AP commercial regions

2. Db2 Warehouse Self-managed
Self-managed, data warehouse on cloud provider managed Kubernetes or OpenShift infrastructure.
↳ Runs on Amazon EKS, Amazon ROSA, Azure AKS, Azure ARO
↳ Available worldwide

3. Db2 Warehouse Cloud Rack
On-premises, OpenShift high performance, data warehouse deployable on Power 10 or X86 infrastructure
↳ Runs on POWER10 or x86 reference architecture

Db2 Warehouse (SaaS) – Gen 3 – Based on Db2 11.5.9

Fully managed / SaaS

Focus on the analytics, we'll take care of the rest

Blazing-fast

Columnar-organized, memory-optimized data warehouse

Scalable & elastic

Independently scale and manage compute & storage

Continuously available

Managed compute, highly available storage, cross-cloud replication

Reliable

Double protection with disaster recovery & self-service backup/restore

Secure

Encrypted at the storage level by the vendor and customer can protect data using their Key Protect keys.

Deploy as fully-managed DBaaS on **IBM Cloud** and **Amazon Web Services**

Moving Forward from PDOA/IIAS – Cloud – Self Managed

IBM Integrated Analytics System (IIAS)

IBM PureData System for Operational Analytics (PDOA)

Migrate to...

1. Db2 Warehouse SaaS
Fully managed, high performance, cloud-elastic data warehouse
↳ Runs on IBM Cloud, AWS
↳ Available in US, EU, UK, JP, AP commercial regions

2. Db2 Warehouse Self-managed
Self-managed, data warehouse on cloud provider managed Kubernetes or OpenShift infrastructure.
↳ Runs on Amazon EKS, Amazon ROSA, Azure AKS, Azure ARO
↳ Available worldwide

3. Db2 Warehouse Cloud Rack
On-premises, OpenShift high performance, data warehouse deployable on Power 10 or X86 infrastructure
↳ Runs on POWER10 or x86 reference architecture

Moving Forward from PDOA/IIAS – Cloud – Self Managed

Db2 Warehouse Software

Container, Operator and Reference Architecture

1. **Container:** Launch a standardized Db2 container across your entire deployment.
2. **Operator:** Easily manage the lifecycle of that container and various database functions.
3. **Reference architecture:** Detailed guide on how to deploy Db2 Data Warehouse on various topologies.

Step-by-step Tutorials for deployment

Db2 on AWS Marketplace

Db2 Warehouse on AWS EKS

Db2 Warehouse on AWS ROSA

Db2 on Azure Marketplace

Db2 Warehouse on Azure AKS

Db2 Warehouse on Azure ARO

Moving Forward from PDOA/IIAS – On-Premises

IBM Integrated Analytics System (IIAS)

IBM PureData System for Operational Analytics (PDOA)

Migrate to...

1. Db2 Warehouse SaaS
Fully managed, high performance, cloud-elastic data warehouse
↳ Runs on IBM Cloud, AWS
↳ Available in US, EU, UK, JP, AP commercial regions

2. Db2 Warehouse Self-managed
Self-managed, data warehouse on cloud provider managed Kubernetes or OpenShift infrastructure.
↳ Runs on Amazon EKS, Amazon ROSA, Azure AKS, Azure ARO
↳ Available worldwide

3. Db2 Warehouse Cloud Rack
On-premises, OpenShift high performance, data warehouse deployable on Power 10 or X86 infrastructure
↳ Runs on POWER10 or x86 reference architecture

Db2 Warehouse *Reference Architecture Solutions*

On-premises offerings that provide an “appliance-like” experience and a smooth upgrade path from PDOA/IIAS (Sailfish) appliances[1].

- Maximizes simplicity, minimizes time to value, and ensures fully supported environments
- *Optimized for performance* and scalability
- **Prescriptive** configuration of IBM hardware, Db2 Warehouse software and Technology Expert Lab services
- IBM built and lab tested
- Includes *Gold-glove Installation* and *Solution Support* Services with a single point of entry.
- X86 & Power-based Solution

[1]As opposed to Appliances where Db2 Warehouse licenses are an integrated part of the solution, in a Reference Architecture Db2 License can be separated from the Solution and moved to another infrastructure.

[1]Clients can also procure the prescribed hardware, software, and Expert Lab services a-la-carte to create a roll-your-own (RYO) or Do-It-Yourself (DIY) self-managed on-premises deployment of a Reference Architecture. However, for RYO/DIY self-managed deployment of a Reference Architecture Solution, IBM will provide only component level installation and support services and **Gold-glove Installation** and ***Solution Support Services*** **will not be available to RYO/DIY Reference Architecture Solutions.**

IBM Power 10 Cloud Rack For Db2 Warehouse

Complete Data Warehouse Solution

Simplify and accelerate your data warehouse deployment in days with our Power 10 based pre-configured solution that includes everything you need to securely analyze your data on your premises – including Servers, Storage, Switches, and Services!

• Simple

- Ease into a Db2 data warehouse with a **pre-configured, all-in-one design** that removes the guess work

• Fast

- Get your on-premises data warehouse up and running in days vs. weeks or months to **start transforming your business faster**

• Efficient

- **Integrate** traditional environments (AIX & IBM i) and new cloud-native applications with Red Hat OpenShift **synergistically**, managed in an on-premises private cloud

• Flexible

- Build the data warehouse needed for today and **easily scale and adapt** as your future needs change

Enable a scalable Power 10 Db2 Data Warehouse that grows as your data increases!

- **Experience:** A balanced resilient architecture for maximum reliability and performance
- **Flexibility:** The Data Warehouse starts small, but grows as you do
- **Efficiency:** E-config enables the full stack pre-selected and configured as a “solution edition” set-up
- **Storage:** IBM FlashSystem adds simplicity and consolidation, capable of delivering more work with fewer drives and lower cost

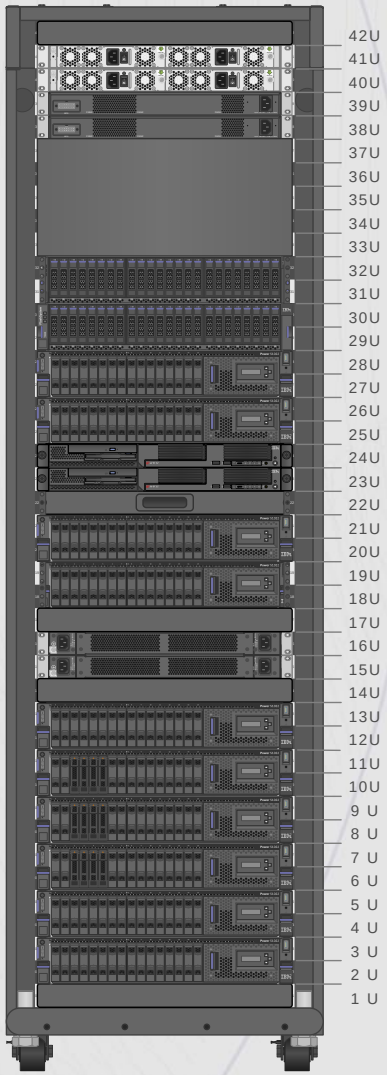

On-Premise HW Stack

Servers:
- OpenShift Management: 3x Power 1022s
- Workers: From 3 to 80 - IBM Power 1022

Storage:
- IBM FlashSystem 7300 for Tier-1
- IBM FlashSystem 5200 for Tier-2

High Speed Networking
- IBM 100 GB Ethernet Switching
- IBM Storage Networking SAN

Management:
- IBM HMC

Warehouse Software Stack

- RHEL OpenShift
- Db2 Warehouse on OpenShift
- IBM PowerVC
- IBM Spectrum Scale Data Management Edition
- PowerVM Enterprise

IBM Services

Pre-sales:
- Size an appropriate configuration for customer’s existing workload.
- Collaboration with business partners.

Install:
- Build the infrastructure hardware stack and pre-load the OCP/SS software stack in IBM MFG
- Gold-glove installation of the infrastructure hardware/software stack at customer’s premises
- Remote installation/customization of Db2 Warehouse Software on the infrastructure at customer’s premises
- Migration Assessment session with the customer

PDOA/IIAS vs. P10 Cloud Rack for Db2 Warehouse
Hardware View

- IBM Integrated Analytics System (IIAS / Sailfish)
  - Power 8 2U Servers
  - FS900 Flash Modules
  - 4x10 Gb bonded network
  - 16 Gb HBAs
  - Db2 Warehouse

- P10 Cloud Rack for Db2 Warehouse
  - Power 10 2U Servers
  - FS7300 Flash Modules
  - 100 Gb bonded network
  - 32 Gb HBAs
  - Modernized Db2 Warehouse on OpenShift

Customer Value

- Substantially improved performance and capacity for the same footprint
- Modernized Db2 Warehouse deployment
- Risk-free ease of migration with options for modernization
- Single-point-of-Contact for all support issues
- Long Solution life and less costly

PDOA/IIAS vs. P10 Cloud Rack for Db2 Warehouse
Multiple Logical Node (MLN) / Hash Partition View

- IBM Integrated Analytic System (IIAS / Sailfish)
  - 8 Data MLNs per Server
  - 3 Power8 Cores per MLN
  - 1.6TB Usable Tier 1 space per MLN
  - 51GB Memory Per MLN

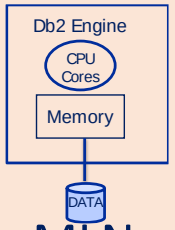

- P10 Cloud Rack for Db2 Warehouse
  - 8 Data MLNs per Server
  - 4 Power10 Cores Per MLN
  - 3TB Usable Tier 1 space per MLN
  - 100GB Memory Per MLN

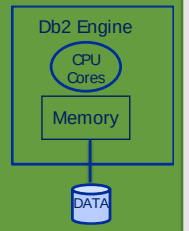

Addition Cloud Rack for Db2 Benefits

- Db2 on RedHat OpenShift embraces Micro Services Multiple Containers vs. One Large Container, **More Granular Upgrades reducing upgrade time**.
- Failover Server per Rack: **Consistent Performance while in Failover Mode**.
- Kubernetes deployment: **Simplifying Initial deployment, Expansion, High Availability, etc.**
- Private Cloud deployment: **Simplifying movement to other Cloud Form Factors**

Cloud Rack for Db2W Base Rack (BR) Configurations

Control Nodes – RHOS Support
Compatible with CPD-S
High Availability Built In

Data (Db2) Worker Nodes
Compute Capacity – Highly Scalable
11-12 Data Nodes per Rack
High Availability Built In

Storage Expansion (Optional)
Storage Capacity
Great for balancing cold vs hot

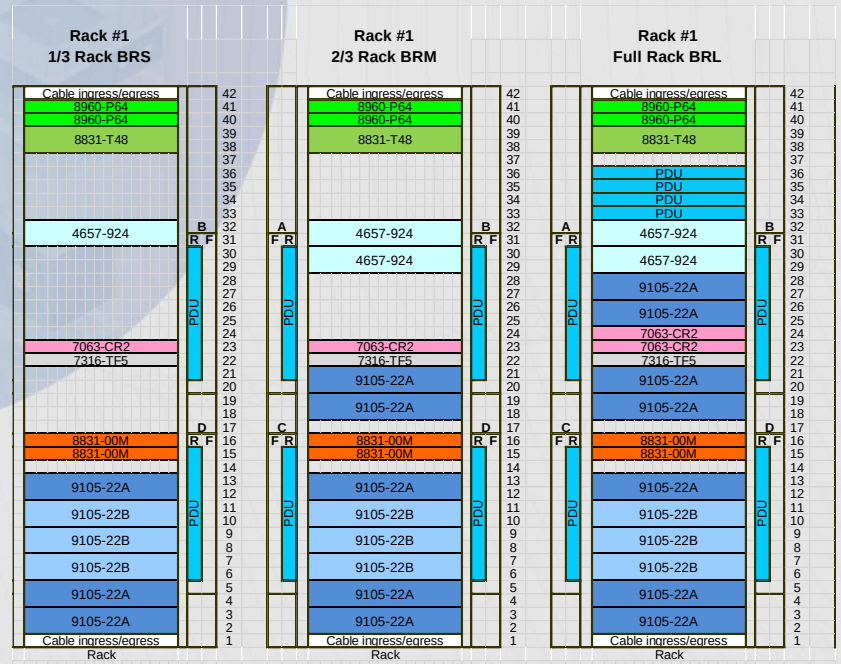

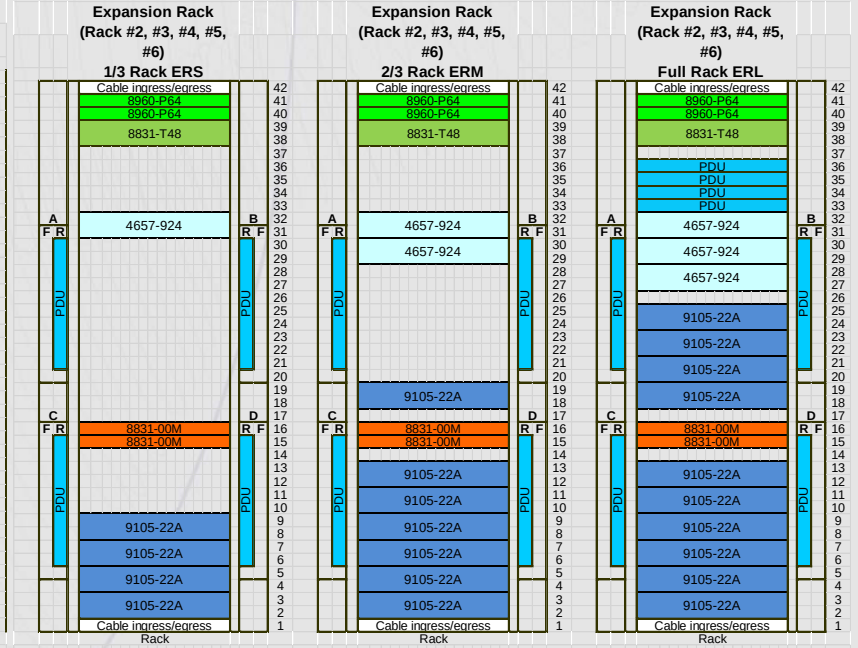

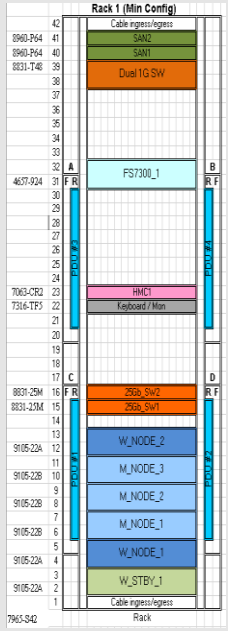

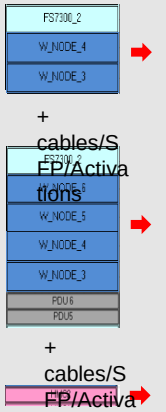

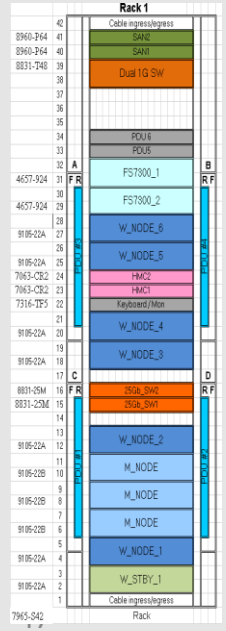

Gold-glove Installation – P10 Reference Architectures

Hardware and OS

- Pre-installation assistance to the client
  - Answer installation/shipping/delivery questions
  - Assist in preparing facilities for the system, e.g., power/cooling compliance, and power/network infrastructure
- On site installation
  - Rack placement
  - Cabling
  - Hardware verification
  - Power on and inspection
  - Upgrades of firmware and OCP/ODF/Spectrum Scale to bring them up to rev
  - network connectivity
  - Sanity check (diagnostics)
  - Stage DB2 SW on system storage
- Post-installation assistance
  - Collect logs

Db2 Software

- Remote installation
  - Install db2u operators
  - Update the configuration of the worker node to match Db2u recommended setting
  - Install db2u entitlement key
  - Create the db2 instance and database (run YAML file)
  - Db2 instances are now Table Ready
  - Hand over the system to the client for ingesting data

- All P10 configurations will be shipped with OCP/Spectrum Scale software installed at the Factory.
- OCP Licenses are included with Db2 Warehouse Cartridges.
- Spectrum Scale licenses are included in eConfig Quote.
- Db2 Warehouse Cartridge Licenses are included in DSW Quote
- OCP/SS/Db2 licenses are enabled during gold-glove installation.

Expertise Connect Technical Account Manager - Ensure Success

• Deliverables

- **Proactive** guidance for preventing issues
- **Remediate** issues
- Achieve application project **milestones**
- **Change management** recommendations
- **Technical** reviews
- **Transformation**
- **Use case** assistance

• Custom Technical Guidance

- Validate your design with **interoperability and security** concerns
- Performance, stability and upgrade readiness **assessments**
- Architecture & health **reviews**
- **Assess skill levels** & identify education options to cover gaps

• Knowledge Sharing

- **Best practices**
- Access to IBM's **product experts**
- Build **your team's competency** on making the most of IBM software
- Information **contextual** to your specific needs
- **Custom tech talks** on key topics

Currently Installed IBM Db2 Warehouse Appliances

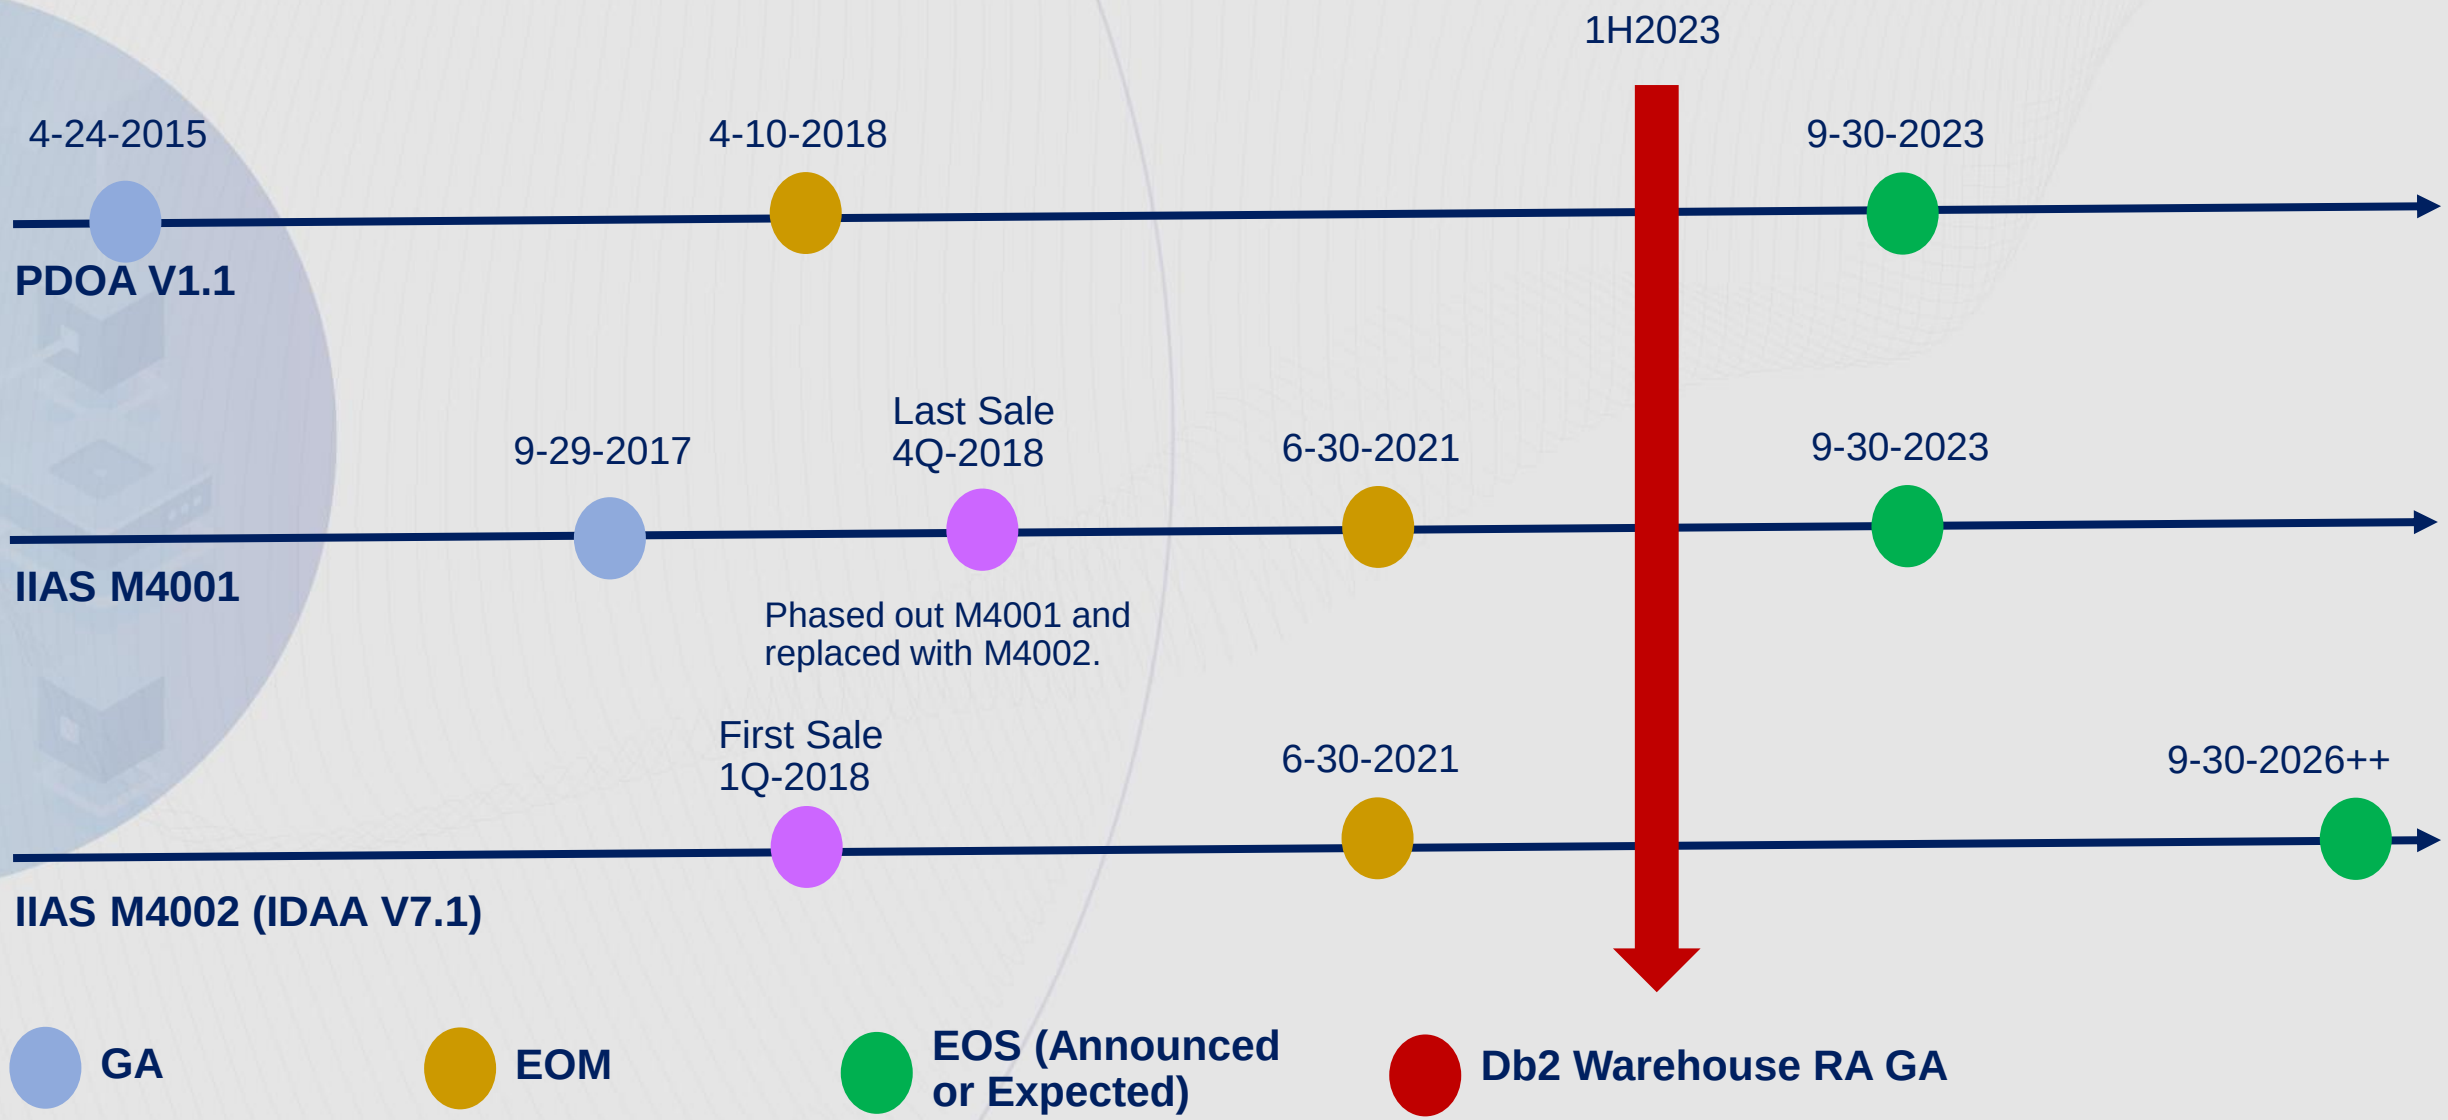

Peak Ahead

Db2 Four big bets for 2024

Continued investment in Db2 on Amazon RDS

Roadmap evolution including new licensing options and other enhancements that make it easy to modernize

Db2 Warehouse Gen3 on IBM Cloud

Fully managed cloud data warehouse featuring Db2 tables on Cloud Object Storage, support for open data formats and watsonx integration

Db2 infused with Generative AI

We’re adding Gen AI capabilities to Db2. Stay tuned.

UX updates for management console

Continued investment to improve the user experience for devs and DBAs

Db2 12

Planned for 2024, Db2 12 will bring significant enhancements to Db2 pureScale, name space separation, generative AI-powered insights, a new AI optimizer and hundreds of other enhancements.

Db2 pureScale improvements

Replacement of TSA with Pacemaker technology for cluster management, leading to significantly faster failure recovery times

Name space separation with TENANT construct

Create a logical separation between one or more database schemas, easily isolating differing sets of tables from each other

AI–powered query optimizer

Allows Db2 to continuously learn from customer's queries and achieve up to 3x query performance improvement over prior version

Db2 infused with Generative AI

We're adding Gen AI capabilities to Db2. Stay tuned.

...and more →

Db2 12

Planned for 2024, Db2 12 will bring significant enhancements to Db2 pureScale, name space separation, generative AI-powered insights, a new AI optimizer and hundreds of other enhancements.

...and more →

- **Improvements to backup performance** by initiating multiple threads to process a single table space
- **Mac M1/M2 driver support** for developers on macOS using Apple Silicon chip
- **Db2 pureScale HADR support for enterprise-grade end-to-end SSL** encryption
- **Online index reorg for Db2 pureScale** allowing index reorg while table remains online/available
- **ADMIN_MOVE_TABLE** performance enhancements
- **Security enhancements** with AUDIT exceptions, Trusted Context and data masking
- Continuing investment in **cloud object storage performance**
- **Schema evolution with DROP and RENAME** support for online schema updates to columnar tables
- **UPDATE and JOIN** performance enhancements for columnar tables
- **Logical backup/restore** experience improvements
- **Recovery time improvements** in the unlikely event of crash
- **Federation enhancements** with support for Snowflake, Oracle 23c and performance improvements

IDUG

IDUG – Important for you

- **IDUG – Active and vital Db2 community**
  - International Db2 Users Group
  - www.idug.org
  - 10,000+ members
  - Db2 LUW and Db2 for z/OS
  - Worldwide

- **Events**
  - India, Australia, Brazil, Mexico, Europe, USA
  - Many regional user groups and events
  - There's likely an active community where you are
  - Face-to-face and virtual events
  - **Bring 3 clients and get 1 free pass**
  - **Buy 1 registration and bring 1 client for free**

- **2024 Events**
  - IDUG India – Bengaluru & Mumbai - February 27-29
  - IDUG NA – Charlotte – June 23-27
  - IDUG Brazil – Sao Paulo – TBD
  - IDUG Mexico – Mexico City – TBD
  - IDUG Australia – Sydney & Canberra – TBD (2025)
  - **IDUG Europe – Valencia – October 27-31**

- **Value & Highlights - Events**
  - 100s of technical sessions
  - IBM Sessions
  - User Sessions
  - Consultant Sessions
  - Db2 Roadmap discussions
  - Product Management & Development Executive briefings
  - Product Management & Development One on One Meetings
  - Special Interest Groups
  - Expert Panels
  - Hands-on Workshops
  - Networking with active community and social events
  - Content for beginners to Db2
  - Badging & Certification Exams
  - Expo hall with demos and experts

- **Value & Highlights – Membership**
  - Education material
  - Presentations
  - Badging & Certification Exams

Survey Questions

Thank You

Speaker: Les King

Company: IBM

Email Address: lking@ca.ibm.com