Db2 V12 Update

Les King
lking@ca.ibm.com

September 2024

Data Server Day – Stockholm, Sweden

Agenda

- Strategy
- Multi-Tenancy
- Availability, Recoverability & Resiliency
- Performance & Scalability
- Multi-Model: Vector Support
- Warehousing
- Db2U
- Db2 Cloud (SaaS)
- Db2 Tools
- Content Design

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

Db2 12 – Delivery Model

- How v12.1 is getting delivered
  - Early access program (EAP) available approximately monthly
  - Early availability in Db2 Warehouse on Cloud – June ✔
  - General Availability all form factors November
- Post GA
  - 12.1 Continuous Special Build Stream (CSB)
    - For Known Issues(KIs aka APARs) available with GA
    - Delivered every 6-8 weeks there after – just like 11.5 stream
  - Continuous delivery of new function delivered in mod packs
  - EAP moves to showcase next mod pack content

Db2 12.1 Supported OS and Platform

Distro	Version	Db2 Server	Db2 Client/DSDRIVER	HA	pureScale
RHEL	9.4	x86-64 System z Power 9,10 LE	x86-64, x86-32 System z (64bit only) Power 9,10 LE (64bit only)	x86-64 System z Power 9,10 LE (Pacemaker)	x86-64 ⭐ System z Power 9,10 LE (Pacemaker)
SUSE	15SP6	x86-64 System z Power 9,10 LE	x86-64, x86-32 System z Power 9,10 LE	x86-64 System z Power 9,10 LE (Pacemaker)	x86-64 ⭐ (Pacemaker)
Ubuntu	22.04 LTS	x86-64 System z Power 9,10 LE	x86-64, x86-32 System z Power 9,10 LE	N/A	N/A
AIX	7.3TL2	64bit Power 9,10 LE	64bit Power 9,10 LE	64bit Power 9,10 LE (TSA)	64bit ⭐ Power 9,10 LE (TSA)
Windows Desktop	11	x86-64	X86-64, x86-32	N/A	N/A
Windows Server	2022	x86-64	X86-64, x86-32	N/A	N/A
Mac	Sonoma ⭐ M1/M2/M3 Arm chip	N/Z	64bit DSDRIVER only	N/A	N/A
UBI Open shift	9 ⭐	x86-64 System z Power LE	N/A	x86-64 System z Power LE	N/A

Multi-Tenancy

Db2 11.5 - Current State

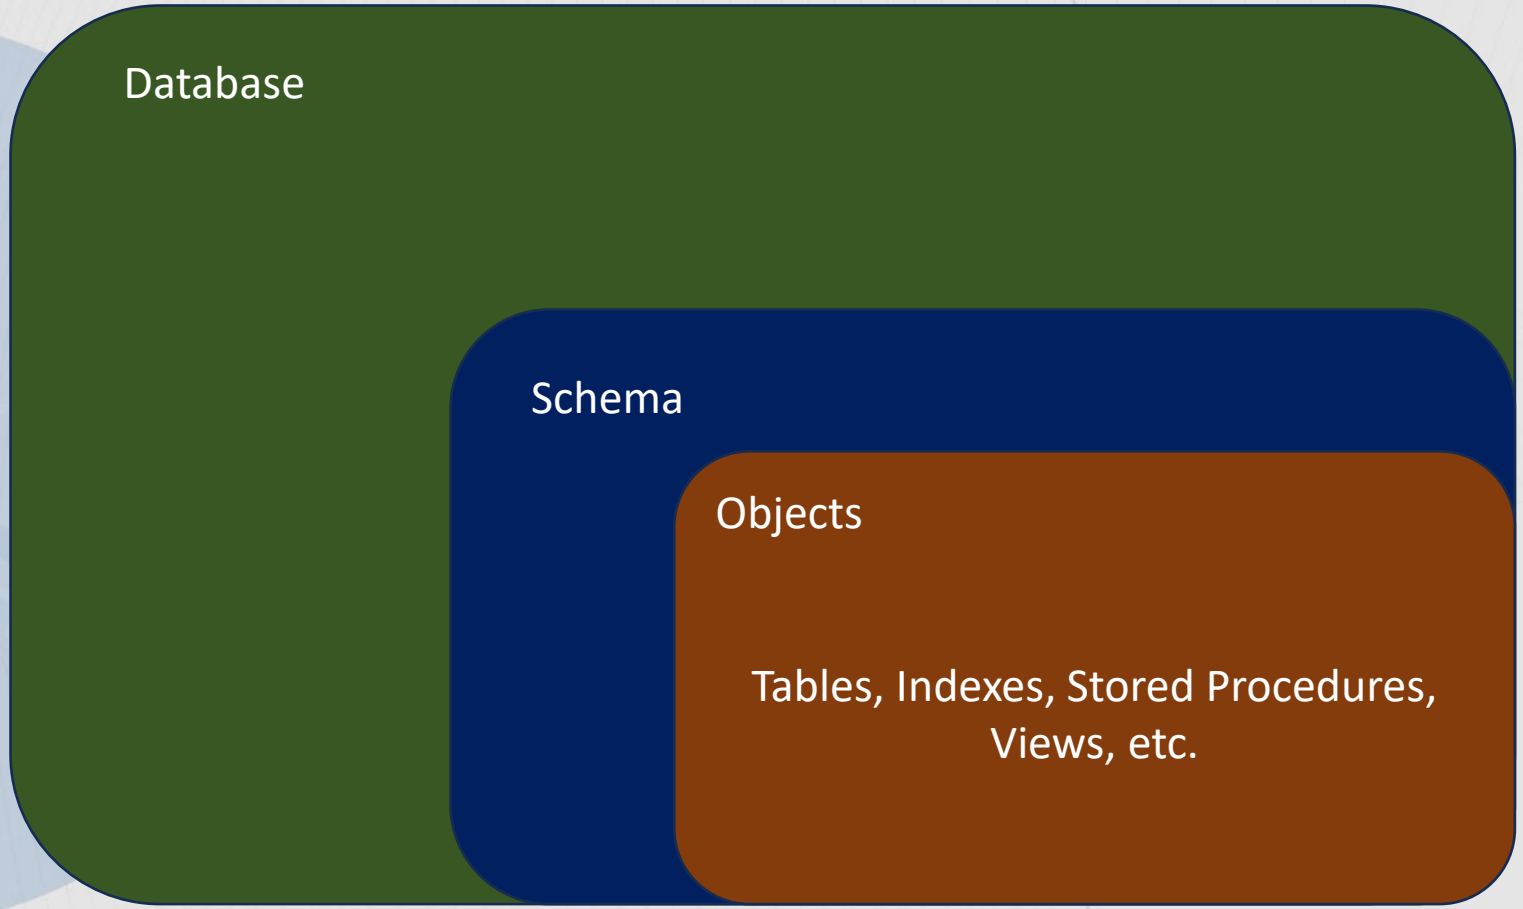

Db2 12.1 - Db2 Tenant

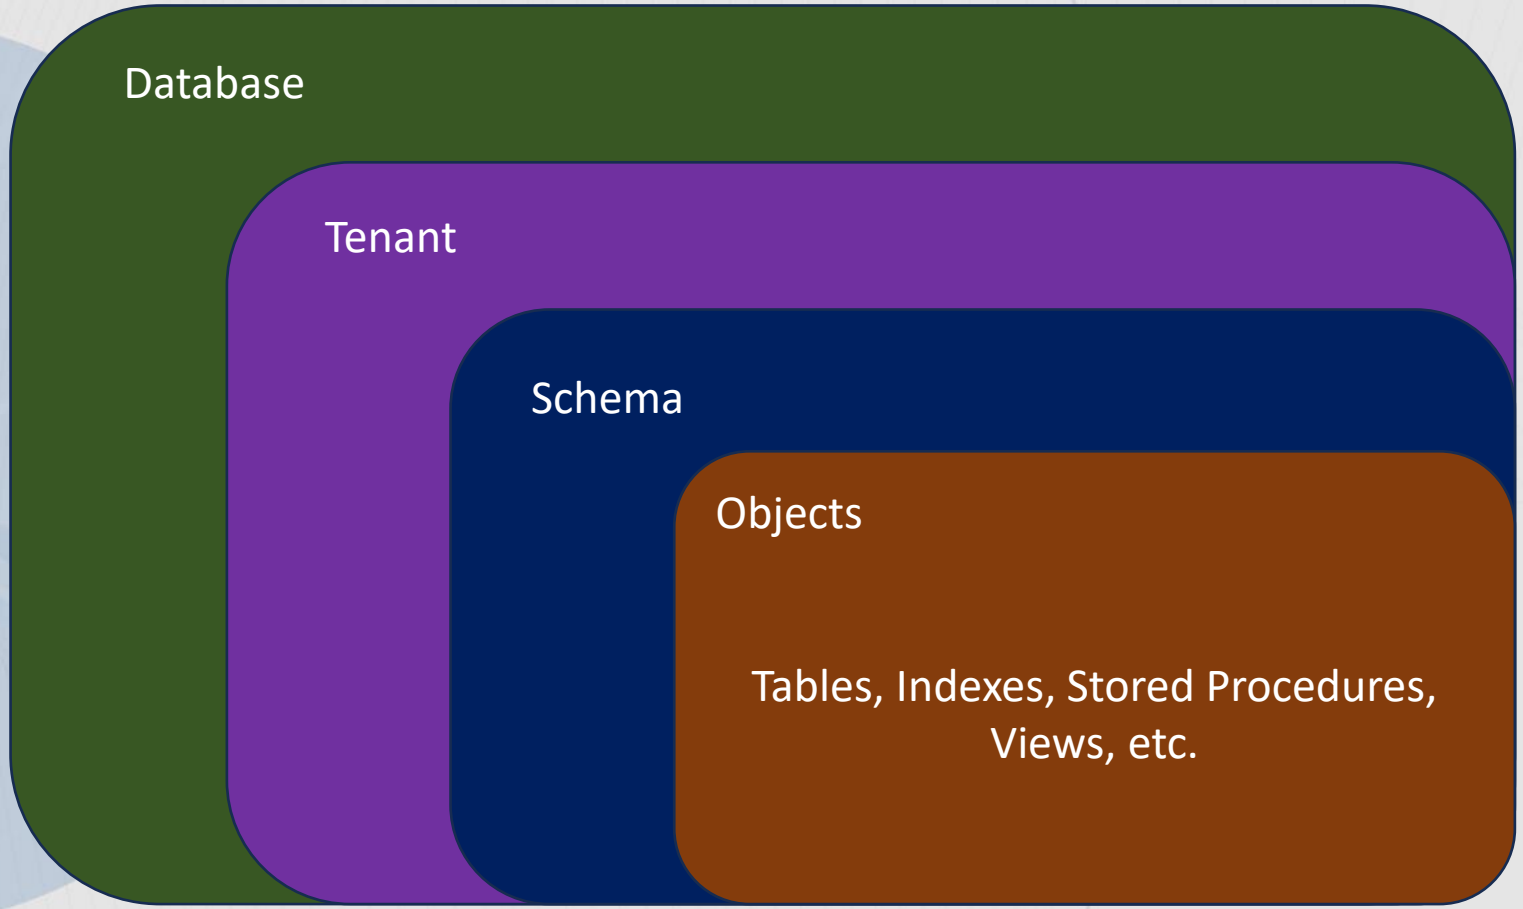

- Change to namespace
- Tenant.Schema.Object

Db2 12.1 – Db2 Tenant – Administration

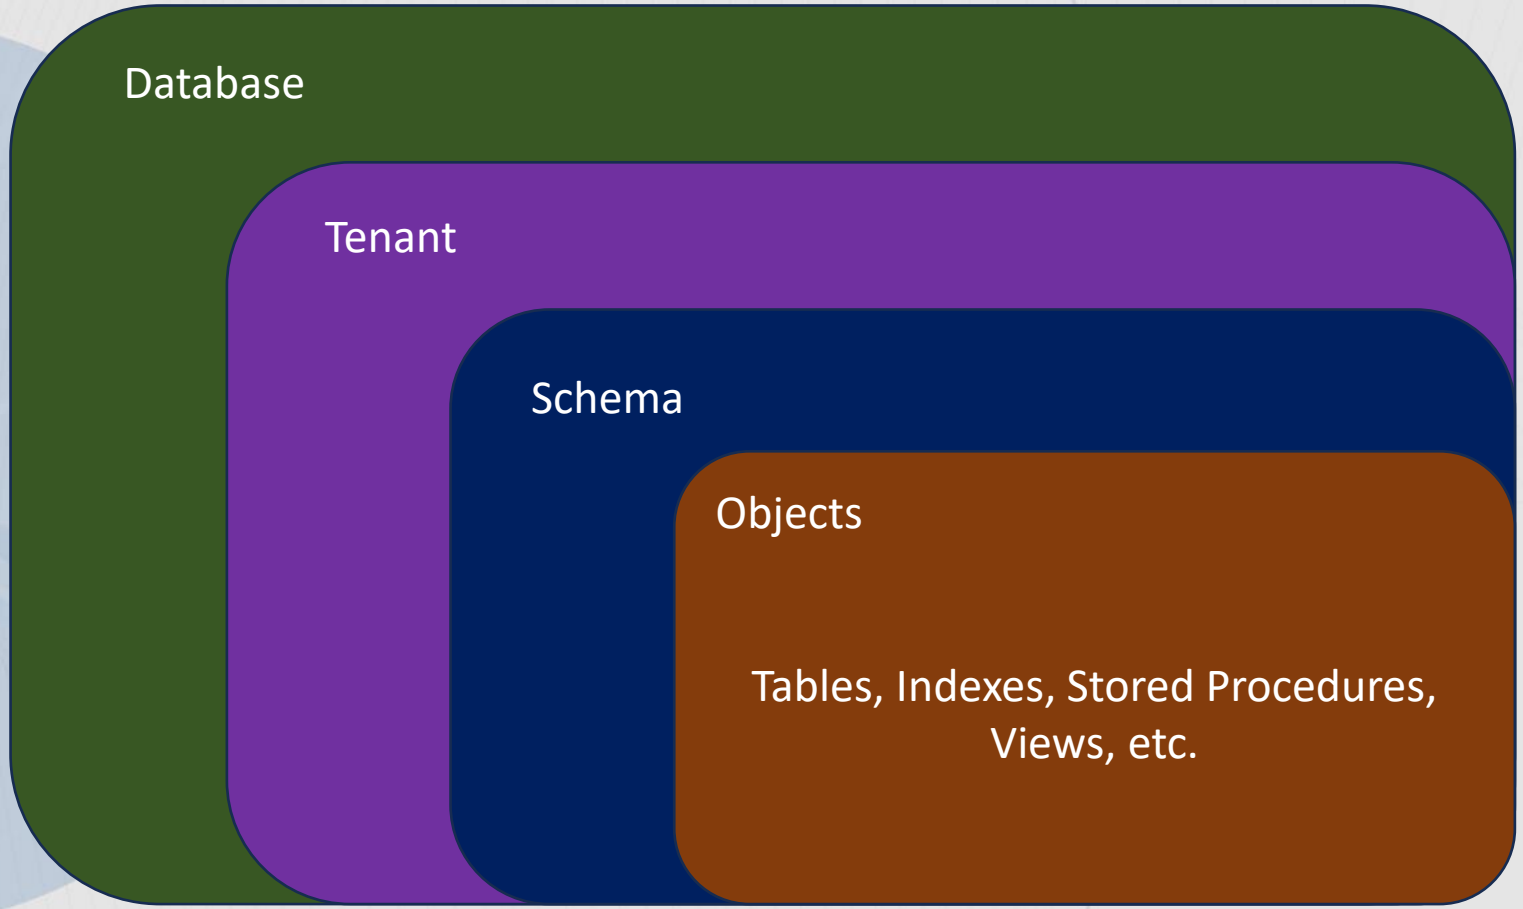

- Tenant-level recoverability
- Tenant-level authorization

Db2 12.1 – Db2 Tenant – Logical Database Support

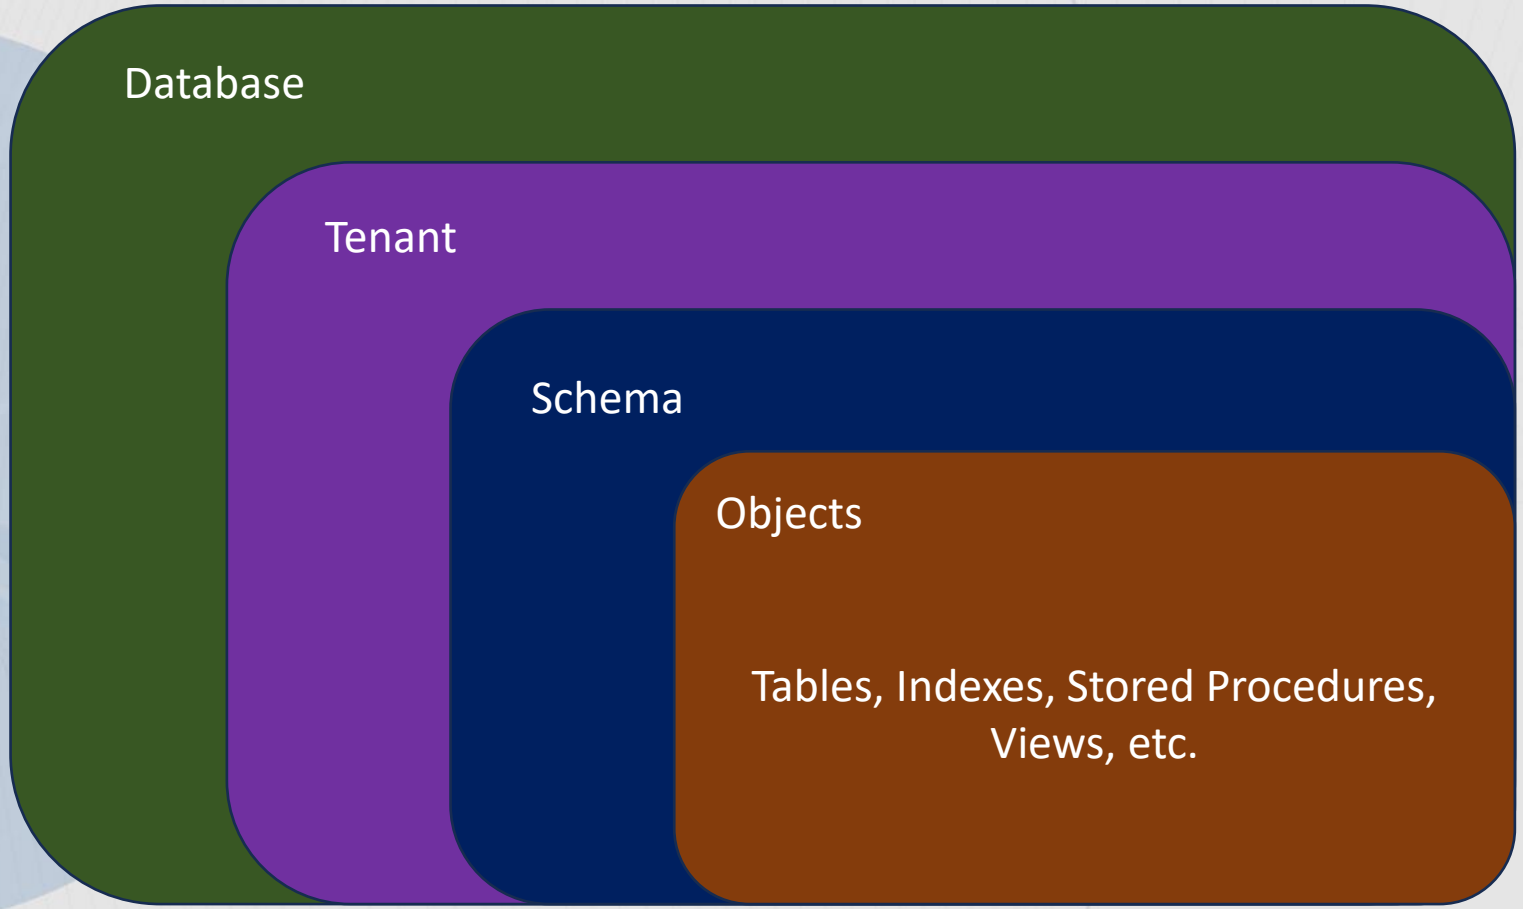

- Physical Database still owns the compute and storage capacity allocations
- Tenant can now be created and be treated as a logical database
- Cross tenant query support will exist

Db2 12.1 – Db2 Tenant – Hosted Application Support

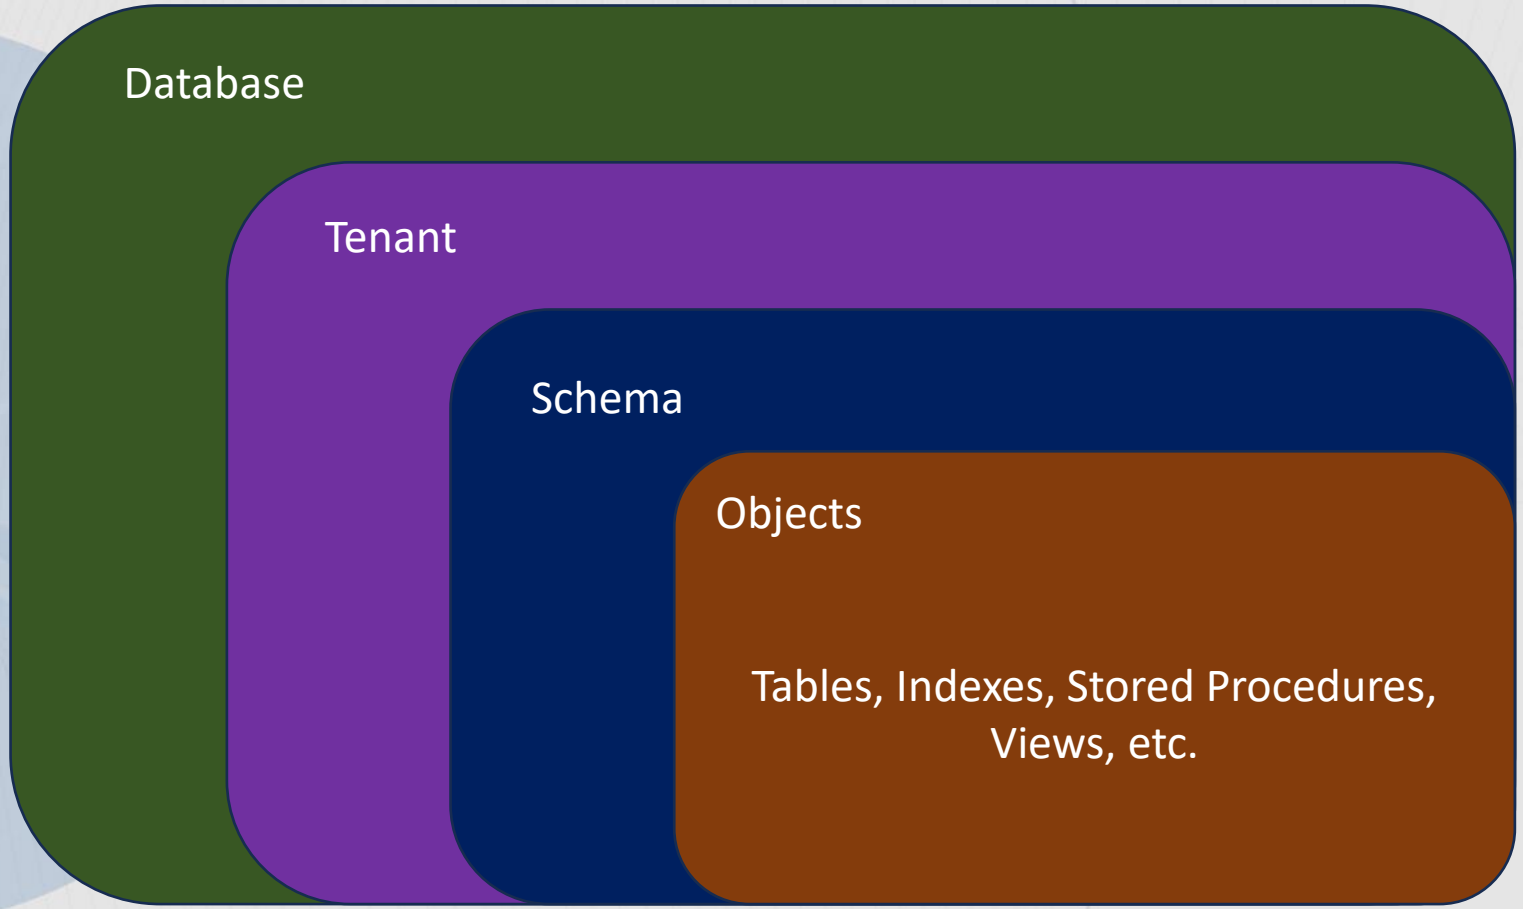

- Hosted application
- Owner of the physical database and database level resources
- Each tenant can have same named schemas and objects
- Each tenant can have their own owner
- Security can prevent cross-tenant access of any kind

Availability, Recoverability & Resiliency

Db2 11.5 – Current State - Pacemaker

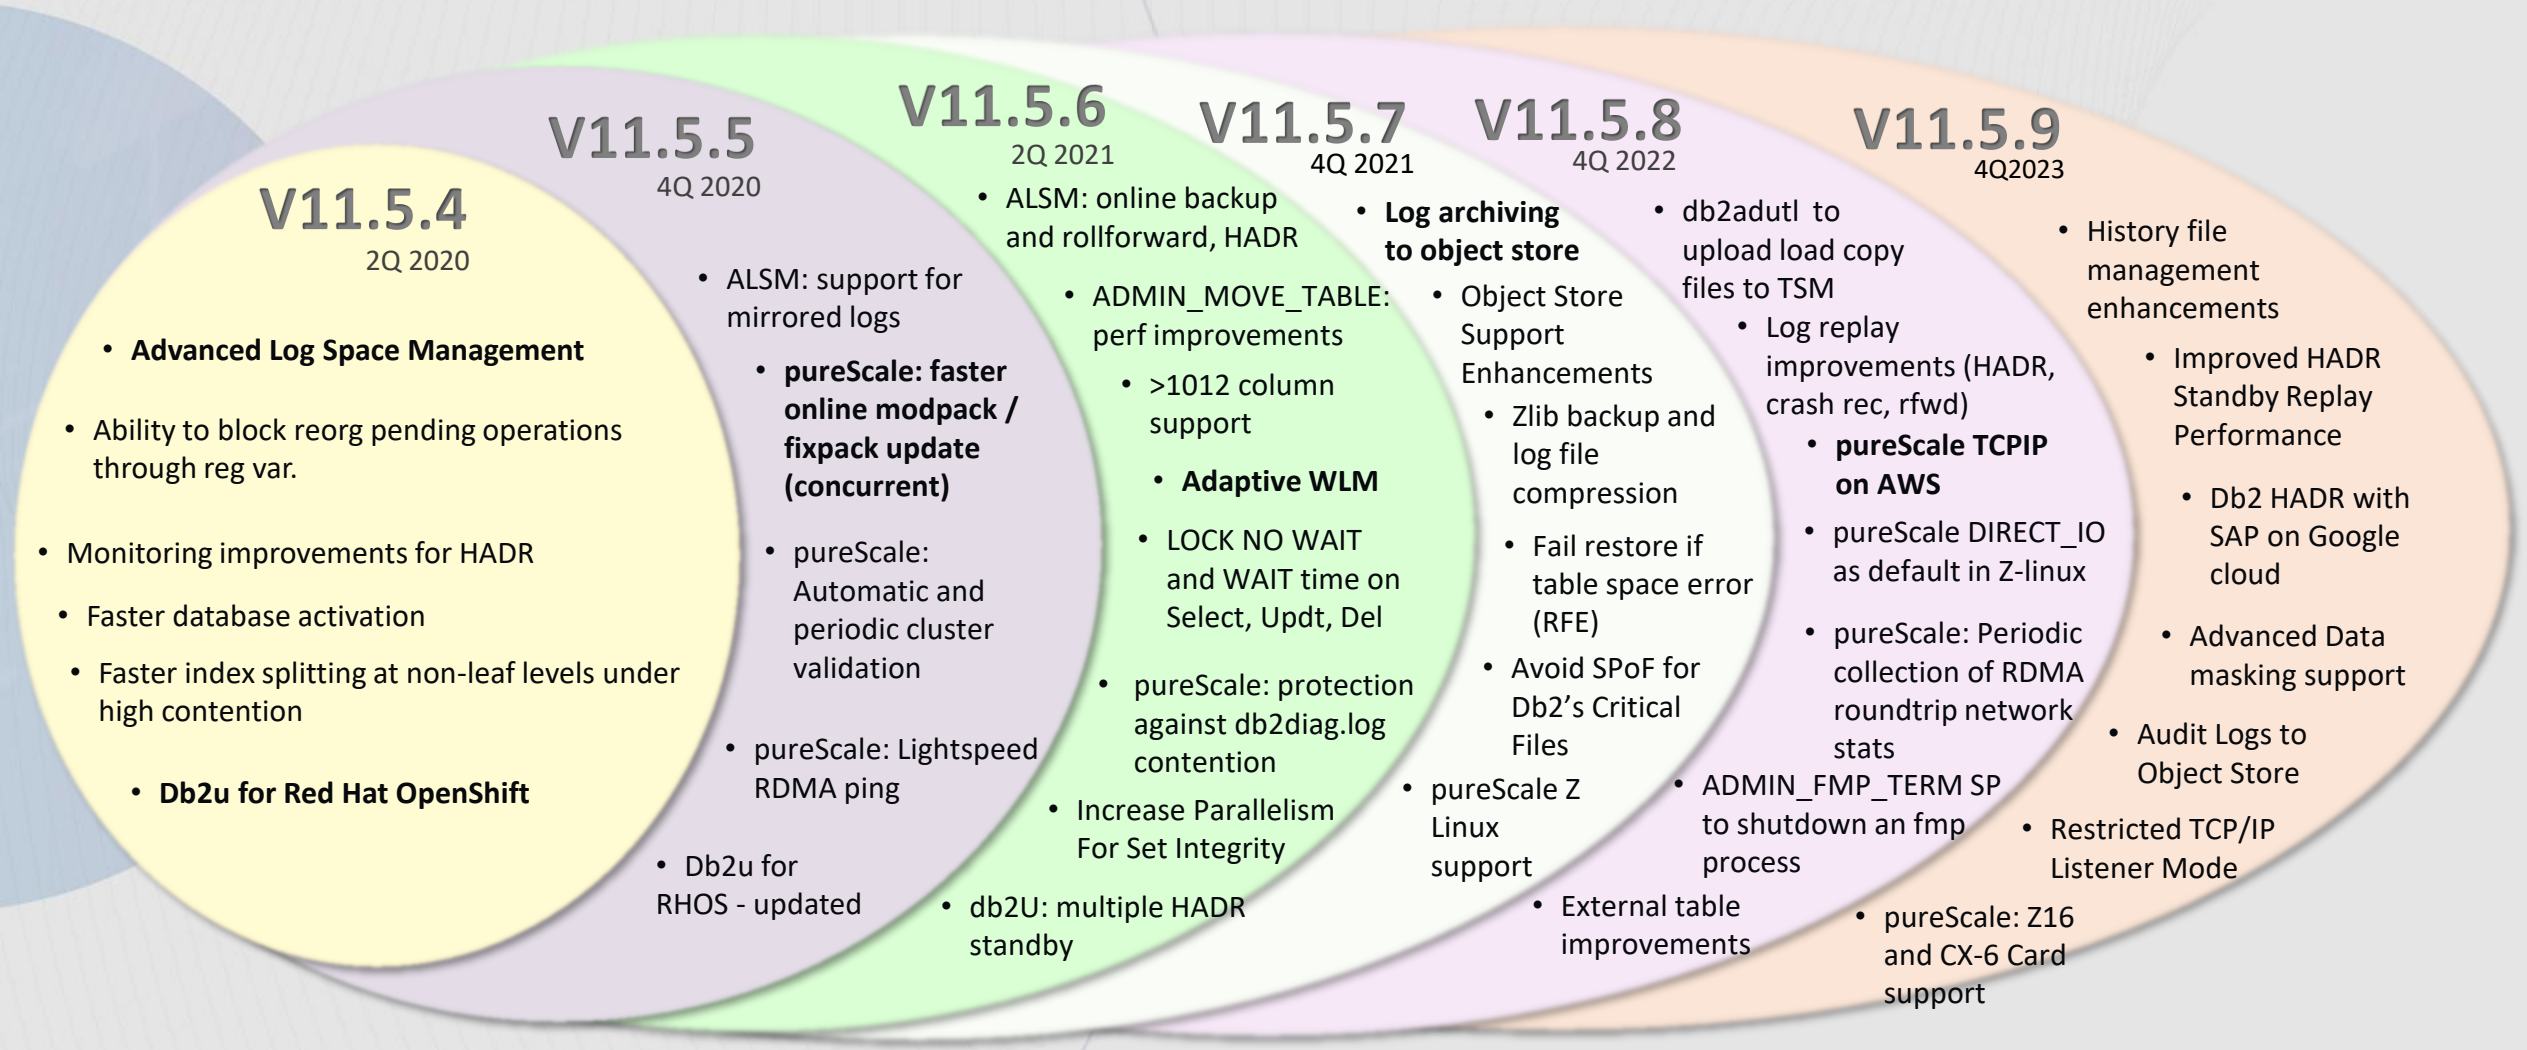

Db2 11.5 – Current State - Pacemaker

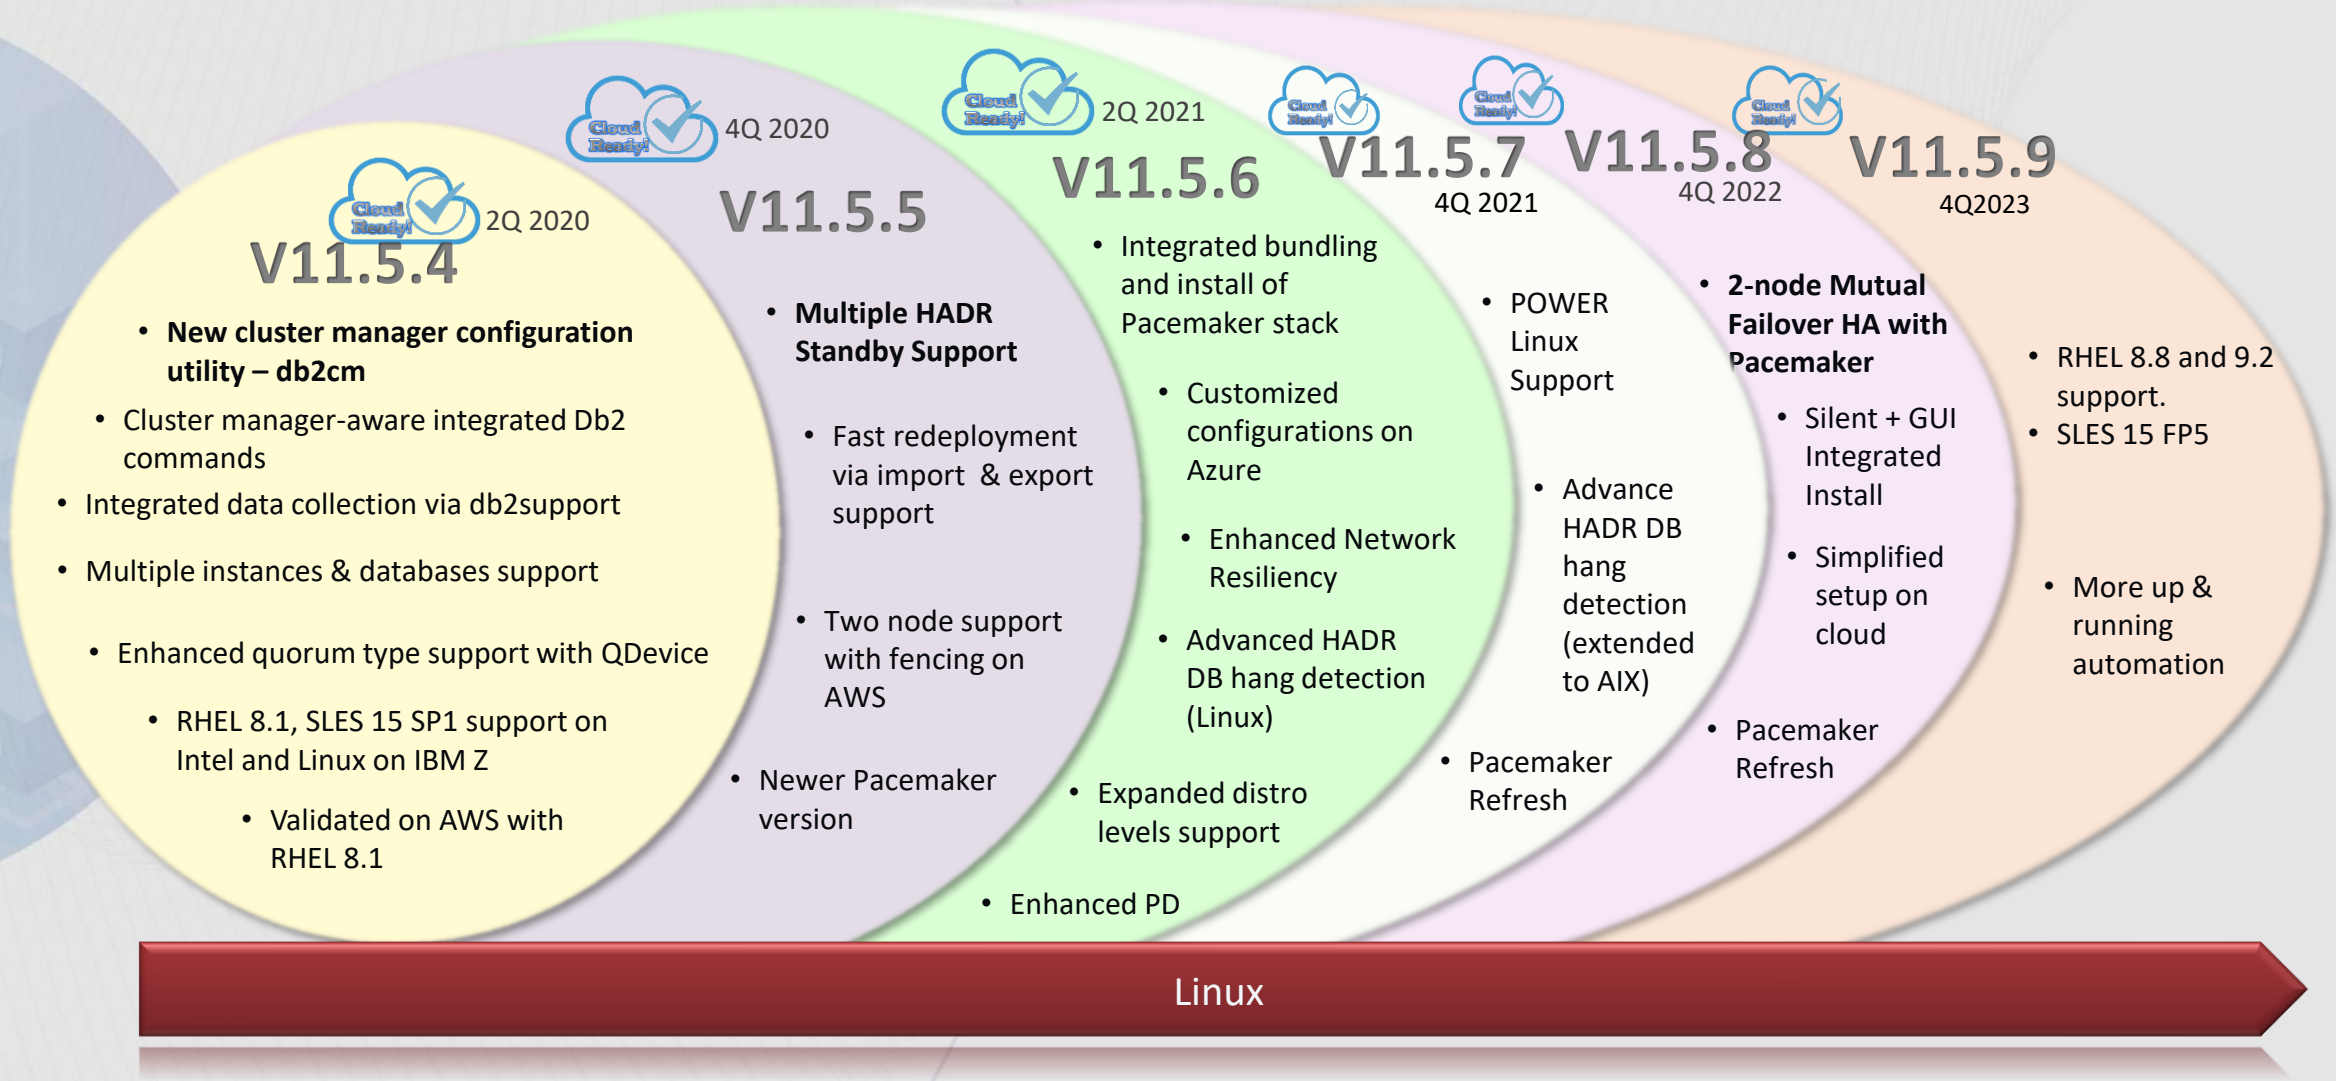

Db2 V12.1 - Availability, Recoverability and Resiliency

Pacemaker for HA Orchestration in Linux

- Replacement of TSA with Pacemaker technology for cluster management for pureScale.
- Replacement of TSA with Pacemaker for DPF HA.

Scalability

- pureScale: Recovery through drop member operations – prereq to support online drop member in 2025.

Communications transport

- Db2 pureScale HADR support for enterprise-grade end-to-end SSL encryption
- Verbs support for RDMA Everywhere in pureScale (SLES on Intel & Z, and AIX)
- Tech Preview: Elastic Fabric Adapter (EFA) support for pureScale on AWS

Online Table Maintenance

- Online Index Reorg (OLIR) in pureScale
- ADMIN_MOVE_TABLE: Performance, usability and functional improvements

Availability

- HADR major version upgrade with RoS
- pureScale - Improve database activation resiliency

BAR and Recovery

- Recovery time improvements in the unlikely event of crash
- LBAR: support storage access alias in place of access keys
- Improvements to backup performance through parallelism at table space level.

Db2 12.1 - Pacemaker as HA Orchestrator

Since V11.5.5

Single DB Partition (EE)
with automated HADR

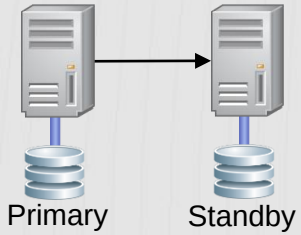

Since V11.5.8

Mutual Failover (a.k.a. Active/Passive)
automated HA with shared storage

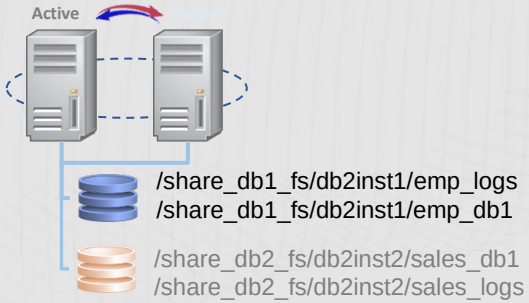

Coming in V12.1

Database Partitioning Feature (DPF)
with automated HA (same site)

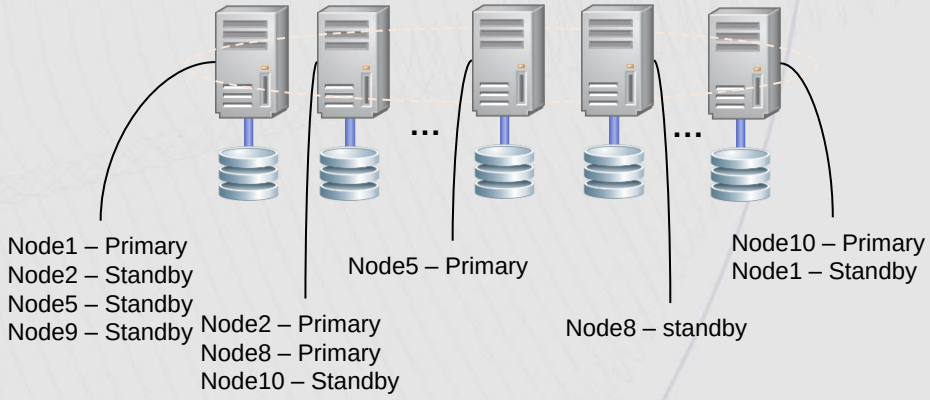

Coming in V12.1

pureScale
Online 24x7x365 with automatic failover

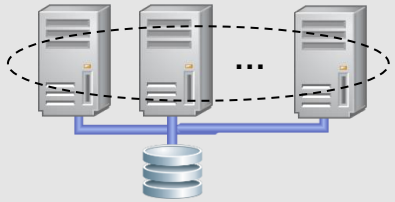

db2cm as management utility

- Unified across all HA configs

Db2 12.1 - pureScale on Public Cloud Roadmap

Nov 2021

Q2 2022 - 11.5.6.0 GA

Self-Managed pS non-RDMA on AWS and TSA <u>with shared storage</u>

- Available in AWS Marketplace
- Light to moderate IUD OLTP workload
- Single AZ. DR via HADR for cross multiple AZ within same region
- Leverage newly released AWS multipath storage

Q4 2022 – V11.5.8.0

Refresh of Q2 GA

- Align with 11.5.8.0 GA
- RHEL support
- Cross region HADR support

Q4 2023 - V11.5.9.0

Update

- Align with 11.5.9.0 GA
- RHEL 9, SLES 15 SP4
- Single push button to setup both HA & DR pureScale clusters

Q4 2024 - V12.1.0.0

Self-Managed pS non-RDMA on AWS <u>with Pacemaker</u>

- Open-source cluster manager

(Technical Preview) Self-Managed pS <u>EFA</u> on AWS <u>with Pacemaker</u>

- Performance comparable to RDMA - suitable for heavy duty OLTP workload
- Fast I/O fencing on AWS (NVMe reservation)

In Progress

Post-V12.1

GA: Self-Managed pS <u>EFA</u> on AWS with Pacemaker
Look into managed service.

GA: Self-Managed pS on Azure with Pacemaker
Look into managed service.

Early Start

Timeline and plan subjected to change

All Clouds

Non-RDMA	RDMA
AWS 2022 Q2	**AWS** 2024/2025
Azure 2025	**Azure** 2025
GCP TBD	**GCP** TBD
IBM POWER VS TBD	**IBM POWER VS** TBD
IBM Intel TBD	**IBM Intel** TBD

Db2 12.1.1+ - pureScale Roadmap

V12 release stream:

The choice for your light, moderate, and most demanding OLTP enterprise workloads on multi-clouds, including both AWS and Azure. The industry-best recovery time and highest availability, fully scalable online (out/in), continuous availability with your day 2 operations, cross-cloud replication utilizing HADR and Read-On Standby, and capable of performing online version upgrade to the next major release.

1. Continuous availability with your day 2 operations – more reorg and AMT enhancements
2. Fully scalable online (out/in) 
3. Multi-cloud: Cross-cloud HADR 
4. Read-on-standby for HADR 
5. Mixed topologies in HADR configurations 
6. Online version upgrade to next major release

Performance & Scalability

Db2 12.1 - Query Optimization Enhancements

- Optimization in the Lakehouse
- AI Optimizer
- Enable optimizer improvements delivered in previous releases or service stream
  - Items were disabled by default, under registry variable control
  - Perform thorough performance testing before enabling
  - All are under DB2_OPTIMIZER_VERSION control
- Update catalogs to support future statistics features
  - Collect per-table partition statistics for range-partitioned tables
    - Eventually allow RUNSTATS on table partitions, and roll in/out of per-partition statistics
  - Collect statistics on all DB partitions for DPF systems
  - Allows features to be delivered in future Db2 12.1 mod paks without requiring catalog upgrade
  - Will roll into 12.1.1+

Db2 11.5 – Current State - Optimization in the Lakehouse

Announcement at THINK

Deliver superior price performance

IBM watsonx.data with **Presto C++** v0.286 and query optimizer on IBM Storage Fusion HCI, tested internally by IBM, was able to deliver **better price performance** compared to Databrick's Photon engine, with equal query runtime at **less than 60% of the cost,** derived from public 100 TB TPC-DS Query benchmarks.*

Presto, the open source Linux Foundation project, is a key engine for watsonx.data. **Presto C++** is the latest in the development of Presto 2.0, the next-generation version of Presto being developed by Meta, IBM and others that run Presto with Velox, an open source C++ native acceleration library designed to be composable across compute engines. IBM has key maintainers in the Velox project, with contributions to the development of Presto 2.0, including the Parquet and Iceberg readers and support for filesystems. **Query optimizer** integrates enterprise-proven query compilation technology coupled with advanced query rewrite and cost-based optimization techniques. In other words, watsonx.data has been enhanced for fast query time performance at optimized costs.

Db2 11.5 – Current State - Optimization in the Lakehouse

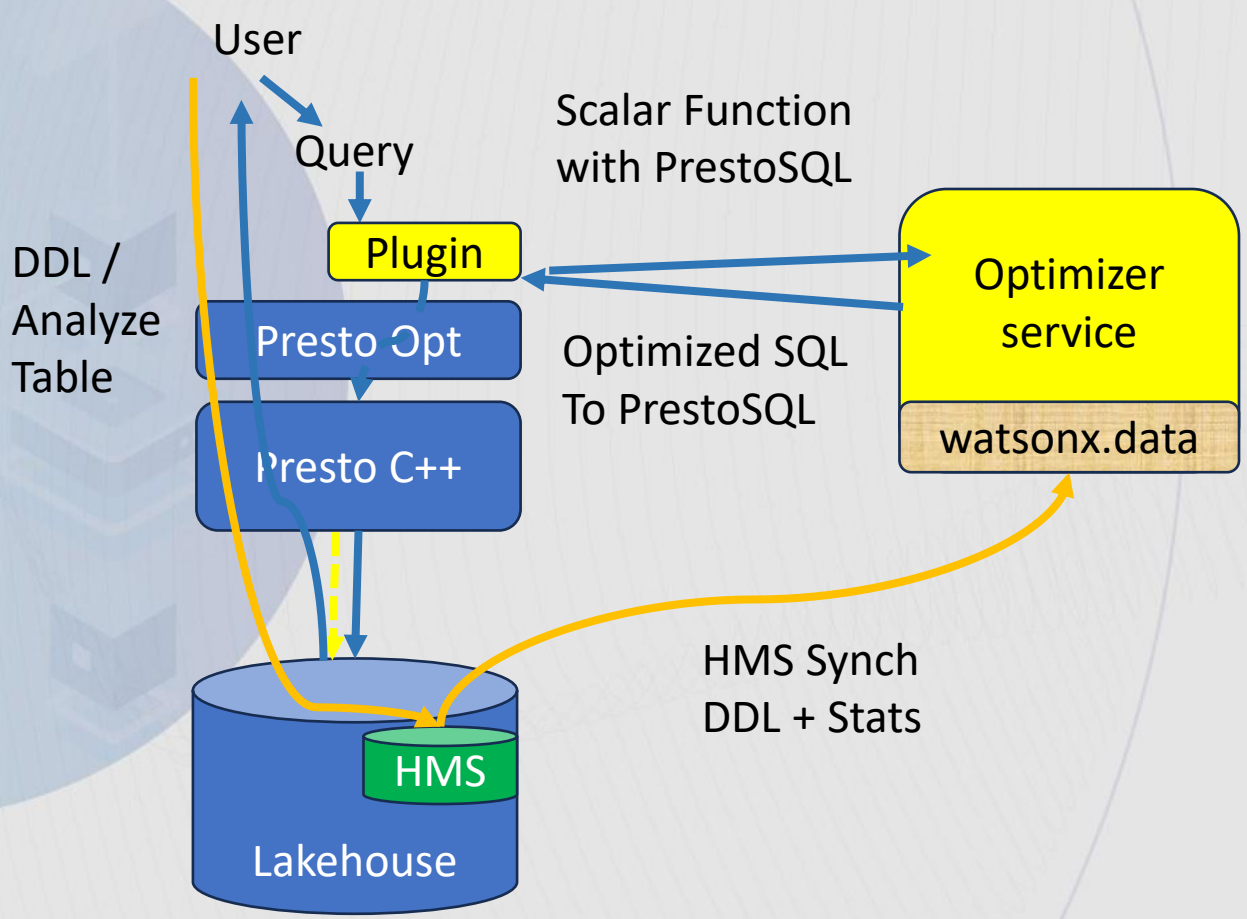

- Translation of Presto SQL to Db2 SQL.
- Complex translation of optimized SQL graph to Presto SQL.
- Some aspects in the Presto Optimizer are suppressed
- DDL / Stats Synch with OptimizerPlus

Db2 12.1 - AI Optimizer

SYSAIMODELS Catalog + SYSCAT Views	Security / Access Control	Audit	Model Policy for control	Activity Logging	DDL for Better Model Management
Better Configuration Management	Appropriate Error Messages	Dependency Management	Increased Accuracy and Number of Columns in the Model	Slightly Larger Model Size 20-30 Kb ➔ 20-60 Kb	Improved Training Time Avg: 60s ➔ 40s
VNext Mod Pack ➔		Some minor enhancement	Explain Support	Join Cardinality using single table models	Optimizer and runtime feedback for model training

Db2 12.1 – AI Optimizer - Infusing AI into Optimization Decisions

Support Card estimation with parameter markers and expressions	Join cardinality beyond single table model use	Index Suggestions	MQT suggestions	Statistical View suggestions	Auto CGS
Registry variable suggestions	Automatic clustering for more effective synopsis	Automatic repartitioning	Query rewrite suggestions	Enhanced Explain Diagnostics	Better Query performance diagnostics in DMC

Open Question: Automatically implement suggestions vs Leave implementation control in user's hands

Likely Outcome: A mixture of the two based on risk and user experience

Vector Support

Db2 12.1 – Vector Support – Background

- A **Vector** is an array of numbers encoding some data.

e.g., [0.6, 0.8, -0.3, 0.7]

- Vector Embeddings: a kind of vectors created with deep learning models
  - Discover hidden relationship among items in a sequence, e.g., words in a sentence
- Vector embeddings are usually high dimensional (100s)
- Foundational to GenAI and RAG applications
  - Retrieval-Augmented Generation (RAG) is the process of optimizing the output of a large language model, so it references an authoritative knowledge base outside of its training data sources before generating a response.

Db2 12.1 – Vector Support - Similarity Search Use Cases

RAGs / Reducing Hallucinations in LLMs	• Appending additional context from vector db to the user prompts and help LLMs generate more accurate content.
Product Recommendations, e.g., Amazon	• Amazon recommending items based on past purchases and browsing history.
Transaction Fraud Detection:	• Banks identifying fraudulent transactions by analyzing patterns in transaction data.

Db2 12.1 – Vector Support - Stage 1

- Goal: enable the storage of AI models encoded in high-dimensional vectors in Db2, next to the data, and do similarity search
- Includes:
  - Vector storage
    - Leveraging existing VARBINARY datatype
  - Built-in Vector Distance Function(s)
    - Cosine distance (most widely used)
    - Perhaps, a few more (e.g., Euclidean)
  - Helper Functions:
    - Vector length
    - Conversion between Db2 data type and Vector string

Warehousing

Db2 12.1 - Data Warehousing Strategy

- ***Cloud Modernization***
- Architecture modernization to exploit functional + economic benefits of public cloud infrastructures and modern deployment models

- ***Core Capabilities***
- Continued enhancements to our core column store capabilities

- ***Lakehouse Integration***
- First class support for open data / open table formats to allow Db2 Warehouse to interoperate in a Lakehouse environment

- ***Gen AI Capabilities***
- New capabilities to support the next generation of applications

Db2 12.1 - Db2 Warehouse (Gen3) on Cloud

1 Available on IBM Cloud

A new IBM Cloud VPC Gen2 platform for new deployments

On par feature set as in AWS, with added support for Flex One models

Next gen cloud storage model and Lakehouse interoperability

2 AWS Improvements

Self serve Private connectivity in AWS

Private Console support in AWS (Connect using AWS Private Link)

Self Serve D/R in AWS. Allows customers to copy Backups to selected regions and restore. Details in a subsequent slide

3 Flexible Updates

Customer's can pick date/time within a window for upgrading their system

Db2 12.1 engine including enhancements to both native cloud object storage and data lake integration

Db2 12.1.1+ - Db2 Warehouse (Gen3) on Cloud

1 Availability on Azure

Db2 Warehouse Gen 3 available as SaaS on Azure Cloud

On par feature set as in AWS, IBM Cloud

Next gen cloud storage model and Lakehouse interoperability

2 Full Separation of Compute + Storage

Functional parity for tables in cloud object storage (indexes, row tables, etc)

Removal of all data dependencies on block storage (eg. catalogs)

More flexible scaling options + data sharing

3 More Seamless Iceberg Table Integration

UPDATE support

Incremental MQTs on Iceberg tables

Optimized performance

Db2 12.1 – Hybrid/Multi Cloud, CDE and Db2 Warehouse Summary

Hybrid & Multi-cloud

Containerization

- Operator driven Db2 native backup and restore
- Db2 WH
  - Vertical Scaling by adjusting resources per pod
  - Horizontal Scaling by scaling out number of pods
  - Support for Native Object Storage
  - Support for Open Data Formats – MPP only
  - Watsonx.data integration – MPP only

- **Tiered cloud storage model** allows tablespace storage to reside on S3 (aka Remote Table spaces)
- **Data Lake Tables** supporting Iceberg, Parquet, AVRO, ORC
- **Columnar MQTs** for Data Lake tables
- **Online Synopsis Table Rebuild**
- Advanced Data masking support
- Audit Logs to Object Store

Columnar Enhancements & Db2 Warehouse

- **AI optimizer** up to 3x query performance improvement over prior version for certain types of queries
- **Improvements to backup performance** by initiating multiple threads to process a single table space
- **ADMIN_MOVE_TABLE** performance enhancements
- **Security enhancements** with AUDIT exceptions, Trusted Context and data masking
- **Faster crash recovery** for columnar tables (up to 6x faster)
- **Logical backup/restore** experience improvements
- **Federation enhancements** with support for Snowflake, Oracle 23c and performance improvements

- **Alter Table with DROP and RENAME column** support for online schema updates to columnar tables
- **UPDATE and JOIN** performance enhancements for columnar tables
- **Alter Table Set / Unset NOT NULL** for columnar tables
- **Alter Table change DECIMAL type** for columnar tables
- **Section Actuals** for columnar tables
- **NLJN** performance enhancements for columnar tables
- **SQL Limits + Illegal Join** support for columnar tables

Db2U

Db2 Universal Container (Db2U)

Modernize your Db2 workloads anywhere !

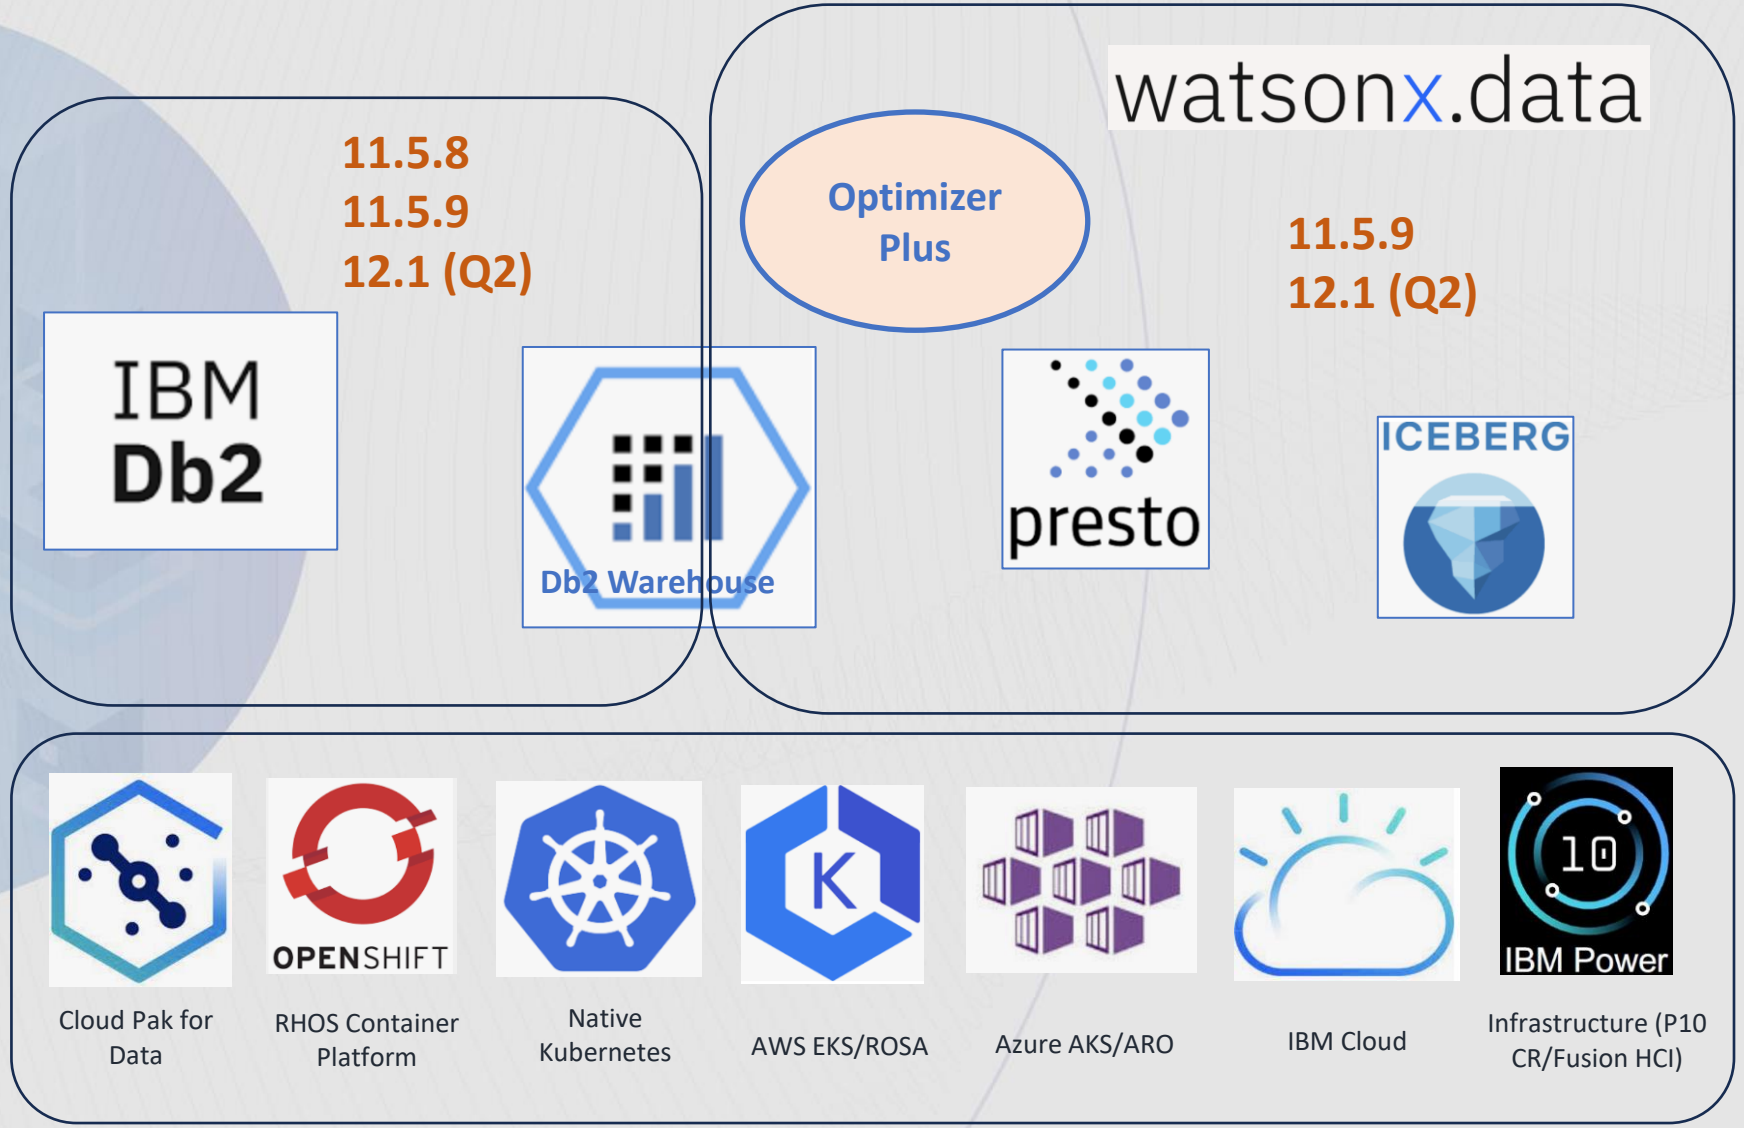

Db2U is the modernization delivery vehicle for:

- Db2 and Db2 Warehouse on CP4D/OpenShift/K8s
- Db2 Warehouse on Cloud (WhoC) Gen 3 on AWS and IBM Cloud VPC Gen 2 (1H 2024)
- Hyper-scalers (AWS, Azure and IBM Cloud)
- Reference Architecture on IBM P10 Cloud Rack and Fusion HCI Infrastructure
- Integrating with watsonx.data ecosystem including Lakehouse (DATALAKE table support)

Db2 11.5 – Current State – Db2U

11.5.8

- Scalability/Availability
  - **Db2U Controller to automatically manage scalability and availability**
  - Scale resources by simply updating Db2uInstance CR resource limits (also in CP4D 4.6.0)
  - Scale from 1 pod to total number of partitions specified during deployment by updating spec.nodes in Db2uInstance CR
- Backup and Restore
  - Driving Db2 Backup and Restore through Kubernetes API **Db2uBAR Custom Resource** (also in CP4D 4.7.0)
- HADR
  - HADR support for OCP and k8s environments. **Automated setup via Db2uHadr Custom Resource**. Semi-automation for >1 site/cluster/namespace.
  - HADR role-aware service that can be used instead of Automatic Client Reroute (ACR) (also in CP4D 4.7.0)
- Audit / Observability
  - Manage Db2 Audit facility via **Db2uAudit Custom Resource** “addOns” (also in CP4D 4.7.0)
  - Cloud-native observability by collecting and streaming key Db2U application logs. Enabled via new Db2uLogging CR. (supports [Cp4D 4.7] Cloud Pak for Data ZEN audit service)

11.5.9

- Db2 Warehouse Native COS support enabled via CR field
- Open Data Format support enabled out-of-the box

OpenShift

2019

AWS **EKS/ROSA** Azure **AKS/ARO**

2022

Infrastructure (P10 Cloud Rack/ Fusion HCI)

2023

[Db2/Db2 Warehouse on AWS EKS]

[Db2/Db2 Warehouse on AWS ROSA]

[Db2/Db2 Warehouse on Azure AKS]

[Db2/Db2 Warehouse on Azure ARO]

[IBM Power Cloud Rack for Db2 Warehouse - Annoucement]

[IBM Power Cloud Rack for Db2 Warehouse – Fact sheet]

Db2 12.1 - Db2U Operator and Container code

Other key enhancements

Db2U architecture alignment

- Deprecate DB2uCluster CR Kind in 1H 2024
- Add support for Db2uInstance CR Kind to CP4D in version 5.0.0
- Remove Db2uCluster CR Kind entirely by 1H 2025 **[TBD]**

Supporting OpenShift default (restricted-v2) SCC for Db2uaaS/non-root installs – 2H:

- Close `MustRunAsRange` gap to align Db2uaaS deployments with OCP default SCC (restricted-v2).
- Expose non-root deployments out-side-of CP4D Db2uaaS to support customers with strict K8s/OCP security posture.

Expanding User management capabilities – Q2:

- Expose all commonly used LDAP/AD-LDAP/ Azure AD configuration settings in `Db2uInstance` CR.
- Unify back-end LDAP/AD-LDAP setup tooling.
- Leverage the LDAP integration capability in the cloud version of the IAM Security plugin to unify authentication mechanisms across CP4D, Standalone and Cloud deployments.

Db2 12.1 - Db2U Operator and Container code
[EXTERNAL] IBM Cloud Rack/Fusion HCI Reference Architecture (RA)

Integrating Spectrum Protect Backup (a.k.a. TSM) – Q4:

- Epic: https://github.ibm.com/DB2/tracker/issues/19539
- Enable TSM integration via Db2U CR “add-On” mechanics. TSM server settings (TSM server, port, etc.) to be exposed via CR spec.
- Adhere to a microservice architecture, where TSM will NOT be installed into the db2u engine container but find a way to load the share libraries.

Contain Pod and PVC placement within a rack - **TBD**:

- Epic: https://github.ibm.com/DB2/tracker/issues/43762
- Introduce a new `placement` field in CR to specify pod scheduling and PV storage allocation.
- In multi-rack configurations of P10 Cloud Rack RA, isolate worker node failures and storage allocation (PVC/PV) to the boundary of a rack to meet the required performance and availability SLAs.
- Consolidate HADR placement requirements in public cloud deployments (topology spread constraints, pod disruption budget) also into the same `placement` field for consistency.

Db2 12.1 – Db2U Operator Maturity Model

Db2U Operator Level 3 roadmap

Day 2 BAR Controller

- [CP4D/Standalone] Native COS BAR support

HADR Controller

- Rolling updates

[NEW] Node Controller

- Detecting cluster node failures: node problem detector
- Checking for node failure events: node healthcheck (NHC) operator
- Tainting failed nodes: Non-graceful node shutdown handling

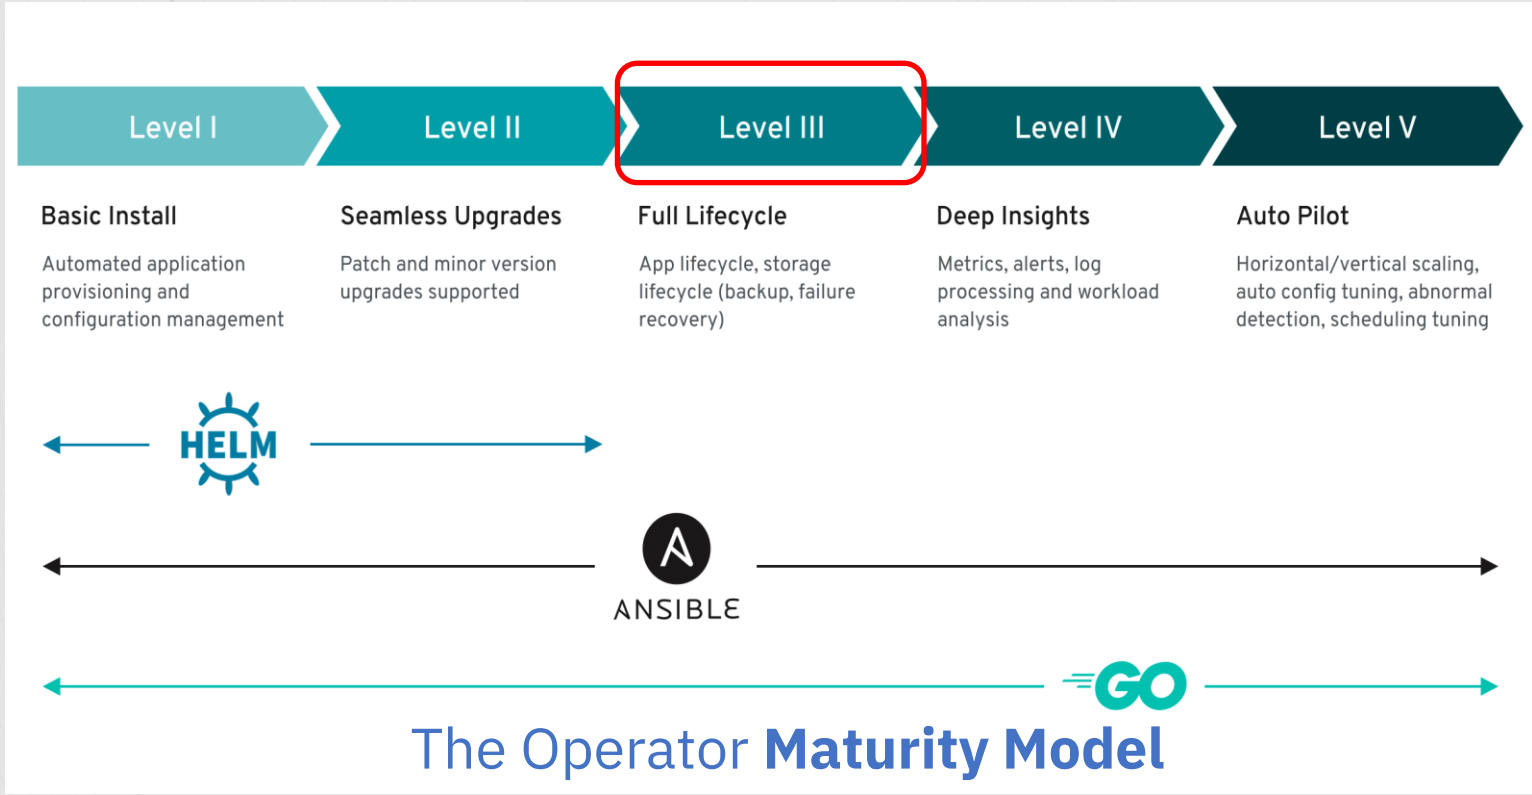

Db2 12.1 – Db2U Operator Maturity Model

Db2U Operator Level 4 roadmap

Day 2 Log Streaming Controller

- Built-in HA (wolverine) HA logs.
- Stream to hyper-scaler endpoints.

Metrics

- Prometheus db2 exporter: Epic 32572
- Explore integration with OpenTelemetry and OpenObserve

Event Generation

- Publishing events from Db2uInstance Controller: EventRecorder

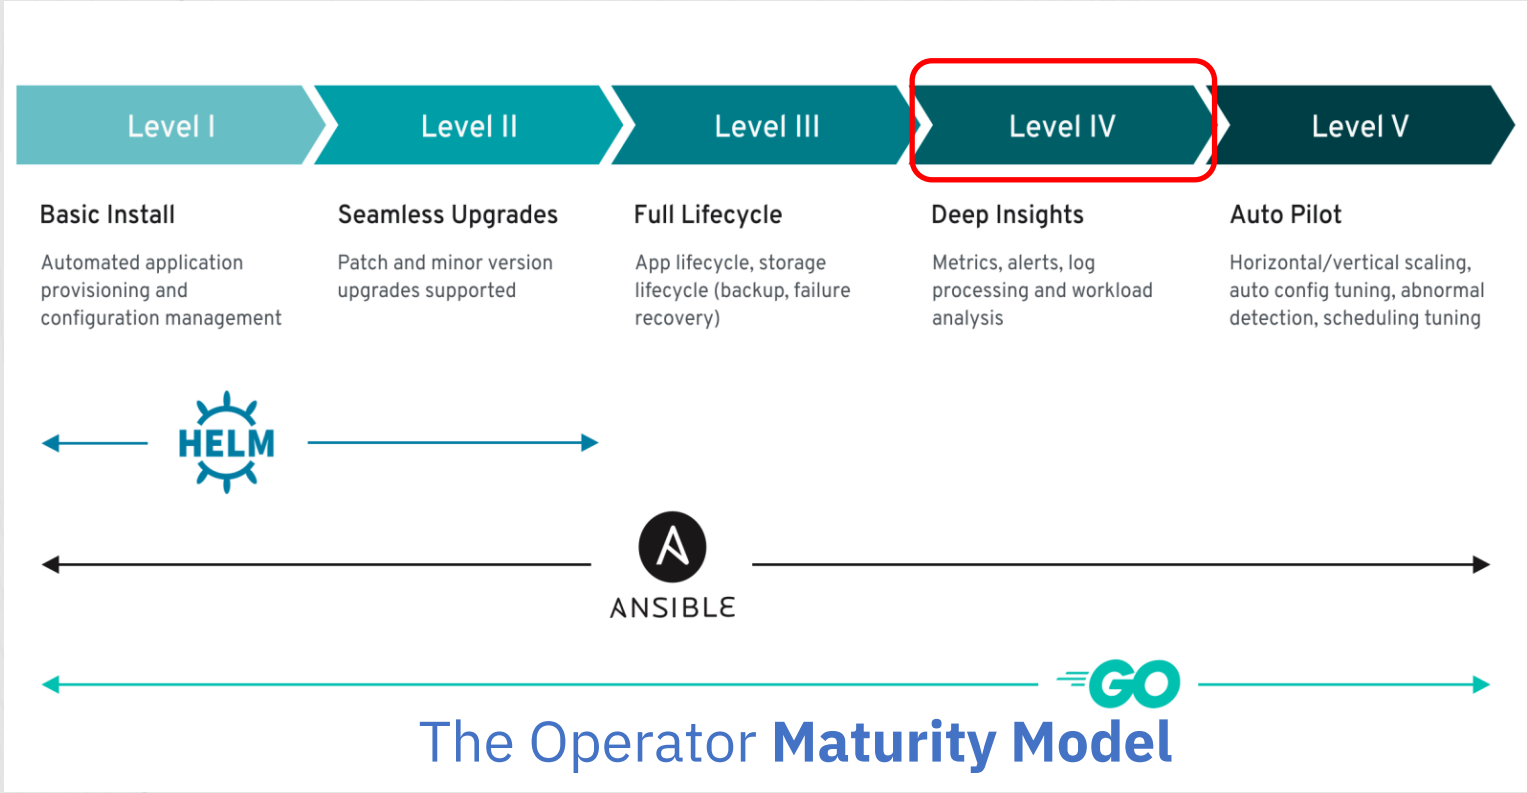

Db2 12.1 – Db2U Operator Maturity Model

Db2U Operator Level 5 roadmap

Vertical Pod Auto-Scaler (VPA)

- Scale up/down (cpu/memory) resources based on workload.
- Supports OLTP/WH SMP/WH MPP.

Horizontal Pod Auto-Scaler (HPA)

- Scale in/out (compute) resources based on workload.
- **[TBD]** Explore for BigSQL first since there already built-in support for online add/drop nodes.
- For Warehouse this will be an offline operation, and scale factor limited to total number of MLNs.

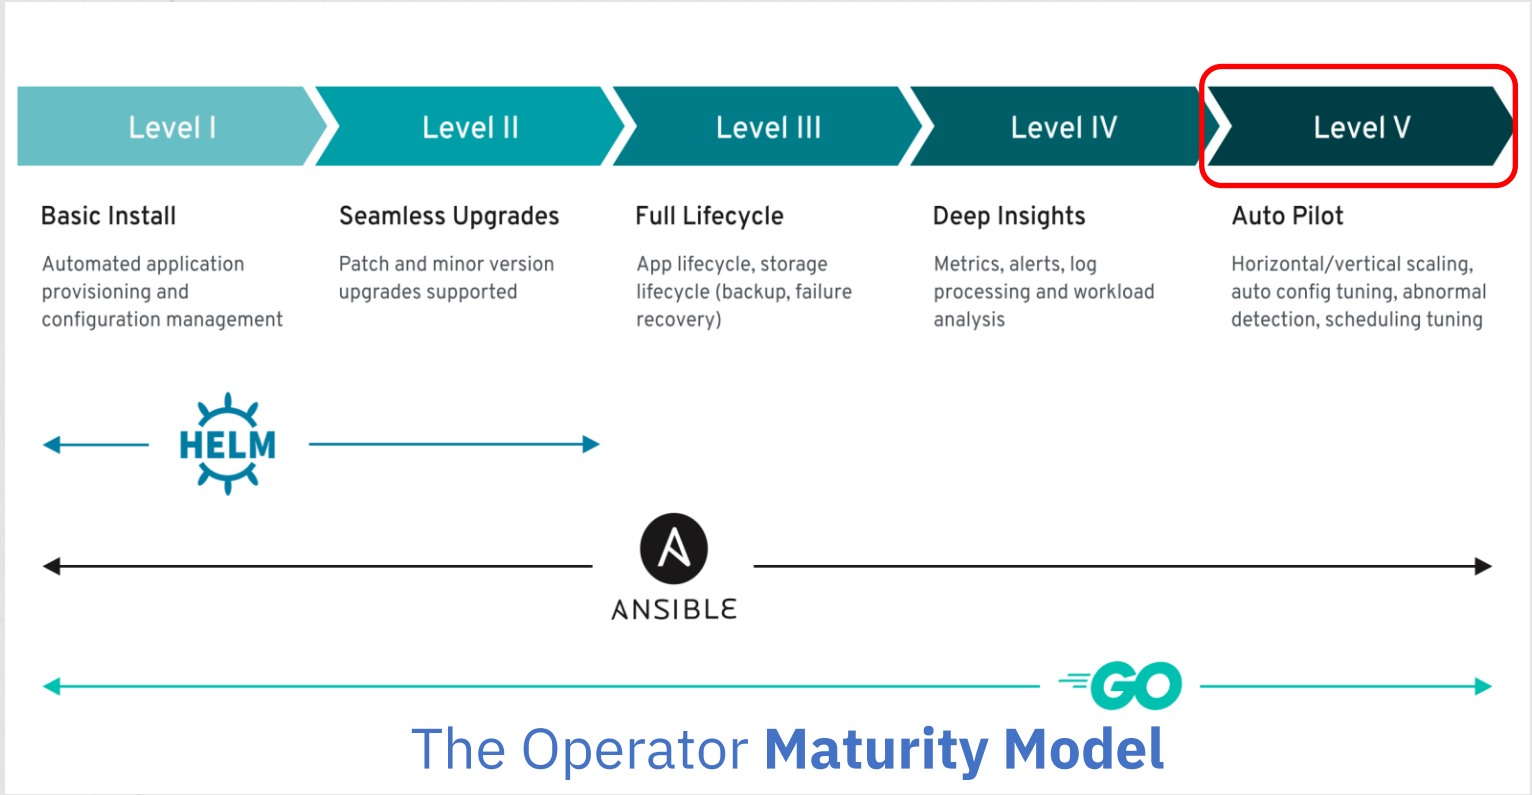

Db2 Cloud (SaaS)

Db2 12.1.1+ - Db2 Warehouse NextGen Platform Architecture

- Delivered on **HDM Common Services Platform**
  - Common Platform for many SaaS services, Db2 on cloud, Db2 Warehouse, DVaaS and now Watsonx.data
  - A similar deployment and management experience for our Cloud Operations team
  - Extensive use of the Kubernetes Operator Framework

- **Control Plane** provides all platform level interactions
  - Flow From IBM Cloud Catalog to Broker for provisioning
  - Billing Flow to send billing to BSS
  - Take a customer request and turn it into a provisioning request
  - Can run on any cloud platform, currently hosted on IBM Cloud

- **Data Plane**
  - Deployed on the native public cloud platform (in this case AWS EKS or IBM IKS) on an **IBM owned account**
  - All customer provisioned components are part of this Data plane
  - Includes the multi-tenant console

- **Runs on Db2u**
  - Db2u containerization provides the underlying Db2 infrastructure (microservice architecture, container images, orchestration – including deployment and scale)
  - Decoupling of compute and storage

Db2 12.1.1+ - Db2 Warehouse NextGen Capabilities

- Scale
  - using the Console or APIs
  - Scale storage or compute independently
  - Operator framework will shutdown everything gracefully , create the new template setting the correct MLN count per pod
  - Brings up everything and orchestrates startup
  - Cluster autoscaler will automatically add nodes or return unused nodes
  - Compute scaling incurs a downtime
  - Storage scaling is completely online
  - Block Storage shrink allowed
- (offline)

- HA
  - Combination of Kubernetes and our own HA component
  - We handle both node and component failures in an automated fashion
  - A very small downtime in case of node failures
  - Console runs multiple replicas, so, no failure if one Console pod goes down

- Backup and Restore
  - Includes both Block Storage Snapshots + Object Storage backups
  - Backup process puts the engine in write suspend mode for the duration of block snapshot only. No halt in processing during this process
  - Restore process will restore both block storage snapshots and restore COS data into the bucket and then startup the pods/containers.
  - Restore Time dependent on COS data that needs to be restored

- Networking
  - Support for both public and private Networking
  - Self Serve
  - Uses standard patterns like AWS Private Link and IBM cloud Virtual Private Endpoint

Db2 Tools

Db2 12.1.1+ - Tools – Migration Strategy

- Db2 Migration Service as the hub for all migration tasks.
- Support all permutations and combinations:
  - Ground to Ground
  - Ground to Cloud
  - Cloud to Ground
  - Cloud to Cloud
- Environments
  - All operating Systems
  - Hardware architectures & big/little endian formats
- Near-Zero Downtime
- Replication/Catchup – when source continues to be active.
- Source to support all possible data sources (target as Db2):

  Oracle, Snowflake, Db2 z/OS, SQL Server, Postgres, Teradata
- Enhanced Dashboard to track progress of migrations

Db2 11.5 – Current State - Tools – HPU & Merge Backup

Now Supports:

- PostgreSQL, Amazon EC2 and remote file system as destinations
- Microsoft Azure as a destination for preparing data for a subsequent upload of it, or for migrating data towards a Microsoft Azure destination, through an appropriate specification of the LOADDEST clause
- Encrypted and compressed backups with the NX842 compression
- Entire database or tablespace for a Db2 for z/OS locally cataloged
- Specification of separate host for load in migration scenario involving a remote Db2 destination
- Accessing to a Db2 Warehouse environment with an SSL connection leveraging the built-in CPU hardware acceleration
- Specification of a general schema for a data migration via INTO TABLES clause with its new WITH SCHEMA option
- Utilizing EXTERNAL output format of Db2, to generate output files ready to be used by a Db2 external table
- Ability of verifying the consistency of the product versions installed across a set of machines by specifying the –check-consistency option in the command line for UNLOAD task
- HPU Current Version is 6.5
- Merge Backup Current Version is 3.5

Db2 12.1.1+ - Tools – HPU & Merge Backup

In Progress:

- Cloud Object Storage as a destination
- Ability to monitor the progress of a running task in terms of amount of data read from database files, backups or the amount of data written to the output etc.
- Ability to unload data from a Db2 environment which authentication method is configured for relying on Kerberos
- Support of Parquet as an output format
- Db2 12.1 Support

Db2 11.5 – Current State - Tools – Db2 Connect & Drivers

Delivered:

- Reactive asynchronous driver for Spring Framework
- Support for IBM Semeru JDK
- .NET driver support for .NET 7 and EF Core 7
- Support for Azure
- Enabling z/OS Batch Runtime with T2zos connectivity for interoperability between COBOL and Java
- Boolean support for Db2 on IBM i

Db2 12.1.1+ - Tools – Db2 Connect & Drivers

In Progress:

- Support for Mac ARM chips M1/M2/M3
- .NET Support for EF Core 8
- .NET Driver for Linux on Z, and Linux on Power
- Support for Visual Studio 2022
- Code First support for Db2 on IBM I
- Db2 12 Support

Db2 11.5 – Current State - Tools – Data Management Console (DMC)

Db2 Warehouse on Cloud Gen 3 introduced:

- Object Storage Support
- Open Data formats – Parquet & ORC
- Copy/move tables to object storage
- Import/export to/from Object Storage
- Import/export to/from Watsonx.Data
- Multi-region backup/restore
- Support for PrivateLink
- Create Datalake Tables in Db2 with link to data in Object Storage
- New Object Explorer

- DMC Standalone
- Cloud Console
- Warehouse on Cloud Console
- Cloud Pak for Data
- IIAS UI/Console
- BigSQL
- Data Virtualization

Db2 12.1.1+ - Tools – Data Management Console (DMC)

In plan:

- DMC to support Power, in providing UI for Power 10 Cloud Rack solution.
- Azure Entra ID support in scope of Multifactor Authentication
- Monitoring and Landing page uplift.

Db2 11.5 – Current State - Tools – Recovery Expert

Recently Delivered:

- Support for PureScale.
- Enhanced pureScale deployment process
- Improved the SSL certificate handling for enhanced security and reliability.
- Current release level 5.5.x

Db2 12.1.1+ - Tools – Recovery Expert

In Progress:

- Enhancing the table data decompression process, including improvements in the dictionary search within logfiles, and the elimination of unnecessary log and backup readings.
- Improving remote Log Analysis support, enabling access to files located on a storage manager.
- Adding integration of Db2 database-level authorities for target databases.
- Simplifying the connection process for a remote computer by replacing the unsecure Remote Execution and Access (RXA) library with Apache MINA
- Allowing datastore upgrade from the Web UI, which previously only supported script upgrades
- Db2 12.1 Support

Db2 11.5 – Current State - Tools – InfoSphere Data Architect (IDA)

Recently Delivered:

- Current version level is 9.2 that:
  - Removed a feature to import/export models
  - Delivered some security fixes.

Db2 12.1.1+ - Tools – InfoSphere Data Architect (IDA)

In plan:

- Support for Java 17
- Db2 12.1 Support
- Product will continue to be in maintenance mode.

Db2 11.5 – Current State - Tools – Data Studio

Recently Delivered:

- Current version 4.1.4
- Db2 z/OS support will end by March 31, 2025. Blog Link

Db2 12.1.1+ - Tools – Data Studio

In plan:

- Product will continue to be in maintenance mode for Distributed Users.
- Support for Java 17
- VS Code Extensions for Db2 z/OS and Distributed to be merged into one.

Content Design

Db2 12.1 - Content Design Enhancements

- Single-sourcing
  - QRep for Db2, Db2 Warehouse, Cloud Pak for Data (CPD)
  - Db2 and Db2 Warehouse on RedHat Openshift and Kubernetes
  - Single-sourcing legacy content for DB2 LUW and Db2 Warehouse (in-progress)
- Documentation for Db2 on Amazon RDS (in-progress)
- Adapting documentation for consumption by Large Language Models (LLMs)
- Improving cross-work with Db2 Support to better address customer issues and reduce support tickets

Db2 V12 Update

Les King
lking@ca.ibm.com
September 2024
Data Server Day – Stockholm, Sweden