IDUG

IDUG

2022 NA **Db2** Tech Conference

Updated (Oct 6/2022)

Modernizing your Db2 containerized footprint with Db2U

Aruna De Silva, IBM (Architect)

Tridex Db2 User Group

Please Note:

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

NOTICE AND DISCLAIMERS :

•U.S. Government Users Restricted Rights — use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.
•Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. This document is distributed “as is” without any warranty, either express or implied. In no event, shall IBM be liable for any damage arising from the use of this information, including but not limited to, loss of data, business interruption, loss of profit or loss of opportunity.
IBM products and services are warranted per the terms and conditions of the agreements under which they are provided.
•IBM products are manufactured from new parts or new and used parts.
In some cases, a product may not be new and may have been previously installed. Regardless, our warranty terms apply.”
•Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.
•Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.
•References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.
•Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.
•It is the customer’s responsibility to ensure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer’s business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer follows any law.

Db2U Overview

What is it ...

- ✓ **Db2 Universal Container** on Kubernetes
- ✓ Standalone & Cloud Pak for Data offerings
- ✓ OLTP and OLAP workloads
- ✓ On RHOS or k8s deployed top of VMs or Bare Metal hardware
- ✓ Cloud-Native User Experience
- ✓ Self-managed offerings on AWS - EKS and ROSA (RHOS)

What's in store ...

- ✓ Self-managed offerings on Azure – AKS and ARO (RHOS)
- ✓ Self-managed offerings on Google Cloud - GKE and Anthos MCM
- ✓ Elastic – Vertical (up/down) or Horizontal (in/out) scaling
- ✓ Cloud-native user-experience for backup, restore, snapshot and other Day-2 operations
- ✓ Db2 Warehouse on Cloud (SaaS) offering on AWS EKS

Roadmap

Db2 Universal Container (db2u)

Recently delivered

Reference architecture: >>
AWS Elastic Kubernetes Service (EKS) AWS Red Hat OpenShift Service (ROSA)
- Db2 Warehouse MPP
- Db2 database single-node

- **Support unique persistent storage volume per MLN**
  Achieve higher performance by leveraging built-in parallelism of Db2 MPP shared-nothing architecture. With the ability to attach a unique storage volume per each database partition, we can support larger scale data warehouses on cloud environments. >>
- **Cloud Storage Support - AWS EFS / EBS**
  Support for storage available through Amazon ROSA. Support includes CPD as well as Standalone. >>
- **Enabling Object Storage Support for Db2 Warehouse** >>
  Enhancements to achieve runtime stability, and improve enablement experience

Q3 2022

Reference architecture:
Azure Kubernetes Service (AKS)
Azure Red Hat OpenShift Service (ARO) >>
- Db2 Warehouse MPP
- Db2 database single-node

AWS Elastic Kubernetes Service (EKS)
AWS Red Hat OpenShift Service (ROSA) >>
- Db2 database HA/DR deployments

- **Support Velero in standalone for both OLTP and Warehouse**
  Integration with Velero to facilitate backup/restore to cloud object storage services. >>
- **Cloud-native approach to Db2 backup-and-restore**
  Operator driven solution to executing Db2 backup-and-restore operations, including Snapshots, thus providing for a better cloud native experience. This support will also integrate scheduling for backup and snapshot operations as well. >>
- **Support leveraging Db2U on Db2 Warehouse on Cloud**
  Support for deploying Db2U via IBM Cloud Control plane (ICD) and Life Cycle Management (upgrade, recovery, etc.) via ICD recipe templates. Day 2 operations via ICD Control Plane Recipes (backup and restore, snapshots, scaling etc.) >>

Q4 2022

Reference architecture: >>
Azure Kubernetes Service (AKS)
Azure Red Hat OpenShift Service (ARO)
- Db2 database HA/DR deployments

- **Release on operatorhub.io**
  Standard for Operator Lifecycle Manager (OLM) releases on non-OpenShift environments (native Kubernetes, Rancher, etc.) for deployment and life-cycle management. >>
- **Improved data warehouse scaling experience for container-based self-service deployments**
  Support horizontal scaling by adding/removing new resources (pods), by simply changing the number of replicas in the Db2u Custom Resource (CR) instance associated with a deployment. Support vertical scaling by incrementing/decrementing the resources given to each Db2u pod, by simply changing the resource limits in the Db2u Custom Resource (CR) instance associated with a deployment. >>

Q1 2023

Reference architecture: >>
Google Kubernetes Service (GKE)
Google Red Hat OpenShift Service
- Db2 Warehouse MPP
- Db2 database single-node
- Db2 database HA/DR deployments

- **Add support for multi-volume per MLN. Stripe Db2 data across multiple volumes** >>
  Helps improve Db2 performance on WH as well as OLTP. Lifting the 1 volume cap (2 TB 4K Block size on AWS EBS), helps with increased storage for Db2 databases
- **A HADR role-aware Kubernetes service for Db2U** >>
  A Kubernetes Nodeport Service for Db2U HADR that is aware of HADR role and routes Db2 client connections to the current Primary POD always. This service can then be exposed via a LoadBalancer or an Ingress service to expose externally. Therefore, from a client standpoint, a single IP is used for catalog/client connection properties

Q2 2023

Reference architecture: >>
Rancher Kubernetes Engine (RKE)

- **Support for Rook-CephFS**
  Provides customers with highly available, multi-use and performant file store. >>

Roadmap

Db2 Universal Container (db2u) for Cloud Pak for Data

Recently delivered

Content underway

Q2 2022

- **Custom Storage Class Assignments Support for RWX and RWO**
  Expands the scope of storage compatibility for all services and enables customers to choose specific class names while installing these services. This objective is linked to support for AWS EFS/EBS. >>
- **Korean language support >>**
- **Support Disruption Free backups**
  Temporary suspend would enable backups without bringing down the pods or the services.>>
- **OCP 4.10 Support**
  Ongoing Cloud Pak commitment to certify even numbered OpenShift EUS versions.>>
- **CPD and Db2 to support VAULT**
  Helps customers to store their credentials to datasources/applications in external vault systems. CPD will retrieve the credentials from these vaults when creating a connection.>>
- **SOD and SCC Compliance**
  Security Context Constraints allows cluster administrators to decide the level of privilege to be associated with any specific workload >>

Q3 2022

- **Individual Service and Service Instance Shutdown and Restart**
  Helps customers reduce the overall costs by 'turning off' the services when not needed, which will reduce the usage of CPU/memory resources and in turn lower the Total Cost of Ownership (TCO). >>

Q4 2022

- **Simplified user experience in enabling add-on products (REST, Graph, Data replication) via CPD UI**
  CPD UI to allow customer specify add-on products, such as REST, Graph, Data replication either during service enablement or as part of the post configuration step. >>

Q1 2023

In Planning

Evolution of Db2 containerization

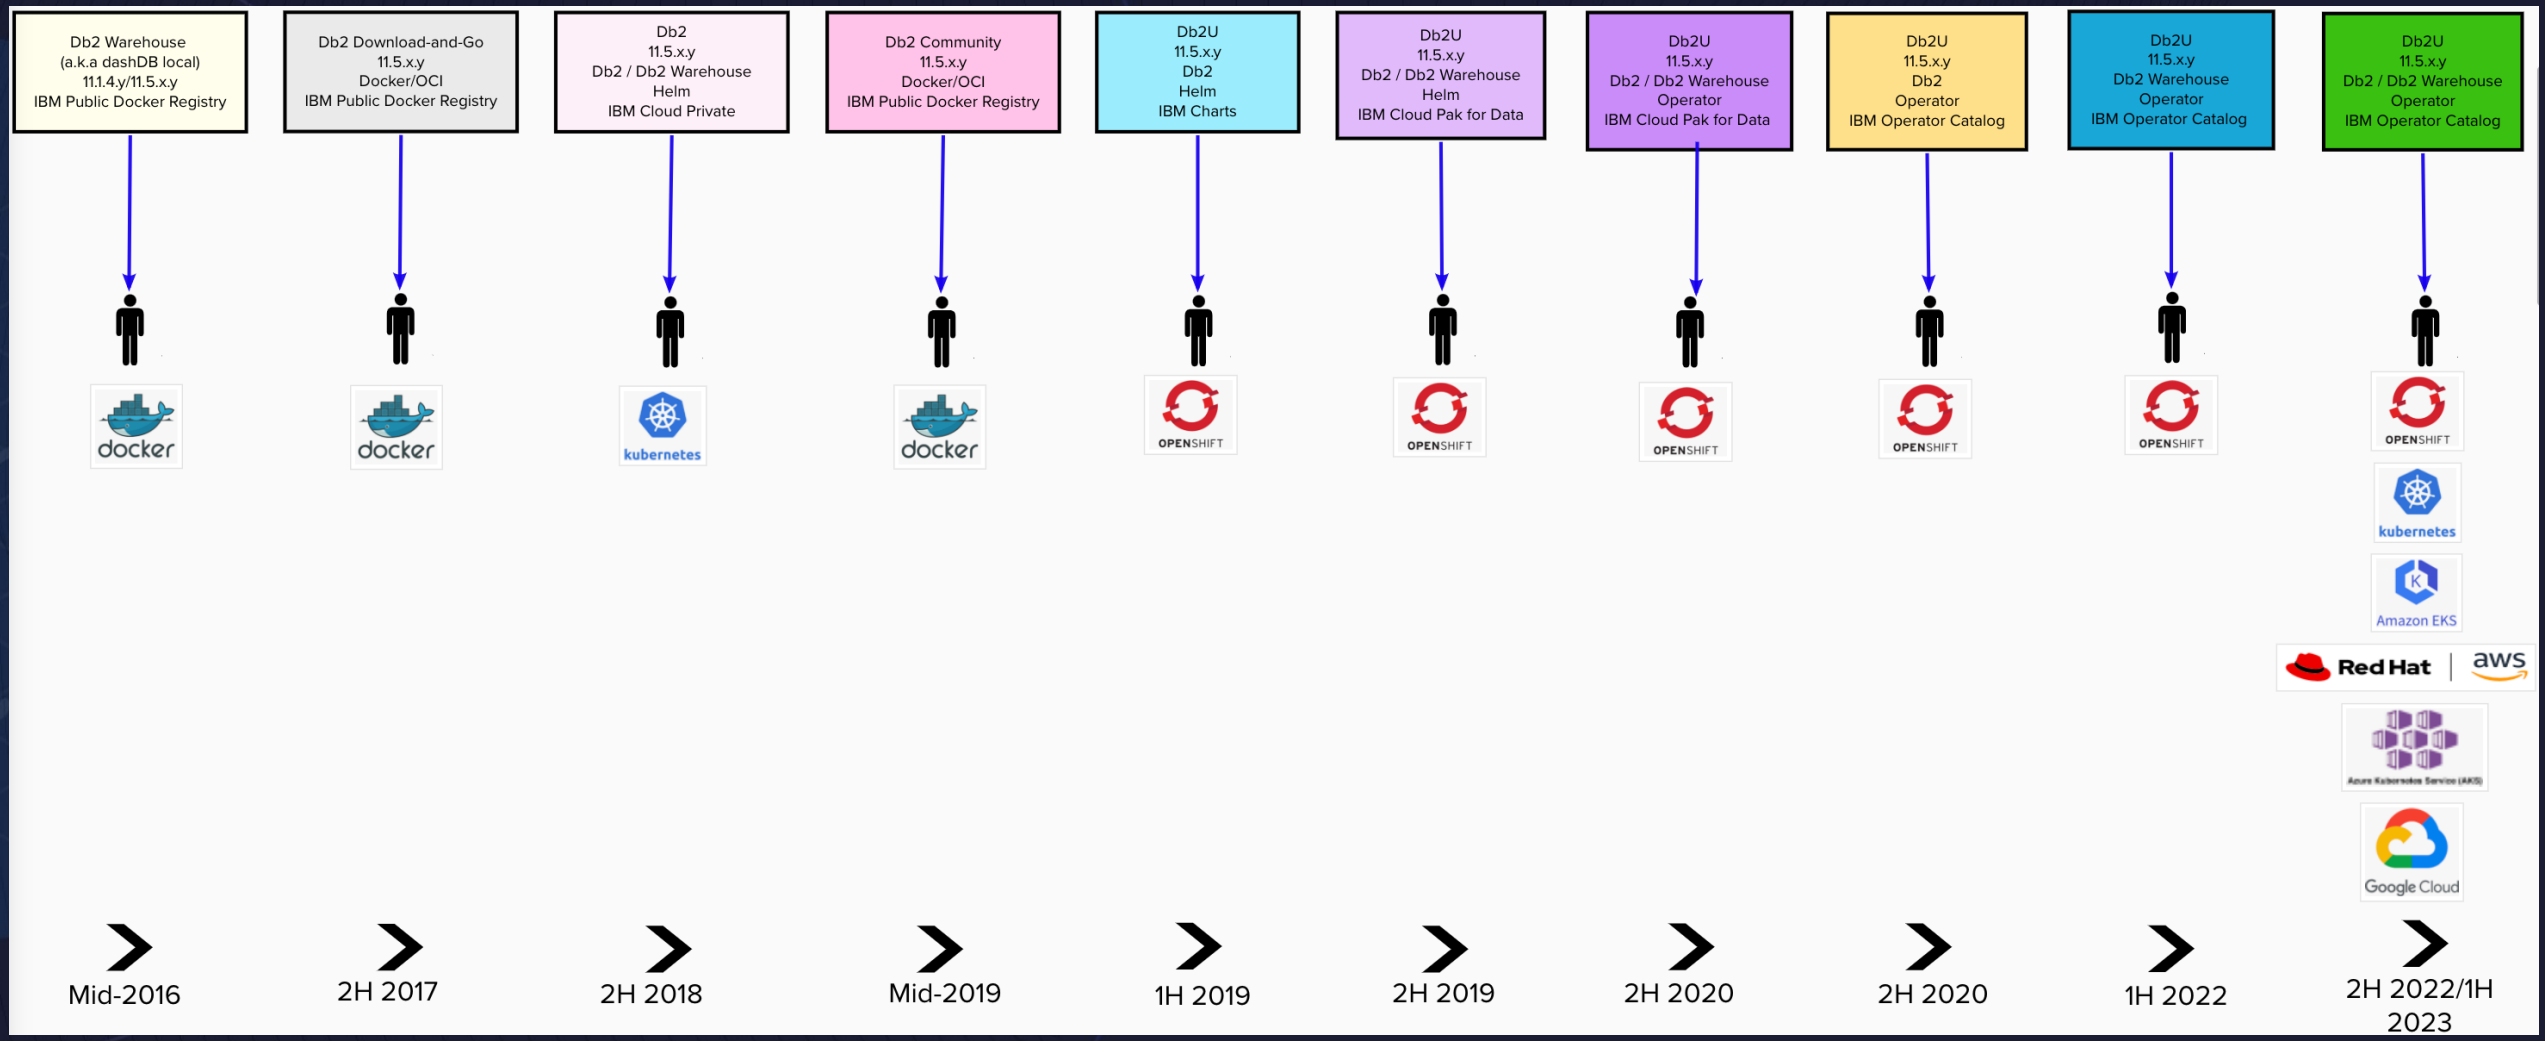

The Db2 Universal (Db2U) Container

- Db2 “Universal” (Db2U) Container – driving Db2 modernization on IBM Cloud Pak for Data, Red Hat OpenShift, Kubernetes and public cloud
  - Microservice architecture
  - Flexible, tailorable form factor – OLTP (Db2), OLAP (Db2 Warehouse)
    - Transaction & data volumes
    - Query patterns & performance requirements
  - Enable pre-built configurations defining the fabric for “infrastructure as code”
    - Portable, secure & certified
    - Ready in minutes
    - Unified environments (consistency through Dev → Test → QA → Prod, etc.)
  - Core to a growing ecosystem of decoupled services

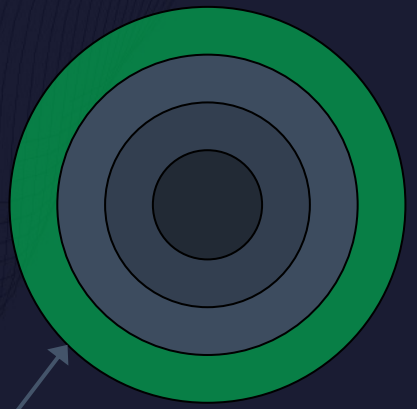

The Db2U Container Hierarchy

An ecosystem of containers built using UBI 8 as the foundation

- ❑ Db2U, BigSQL, Data Virtualization (DV)
- ❑ Add-ons: REST, Graph, Replication (Q-rep)
- ❑ Internal: FVT, Storage Certifier, Release Certifier

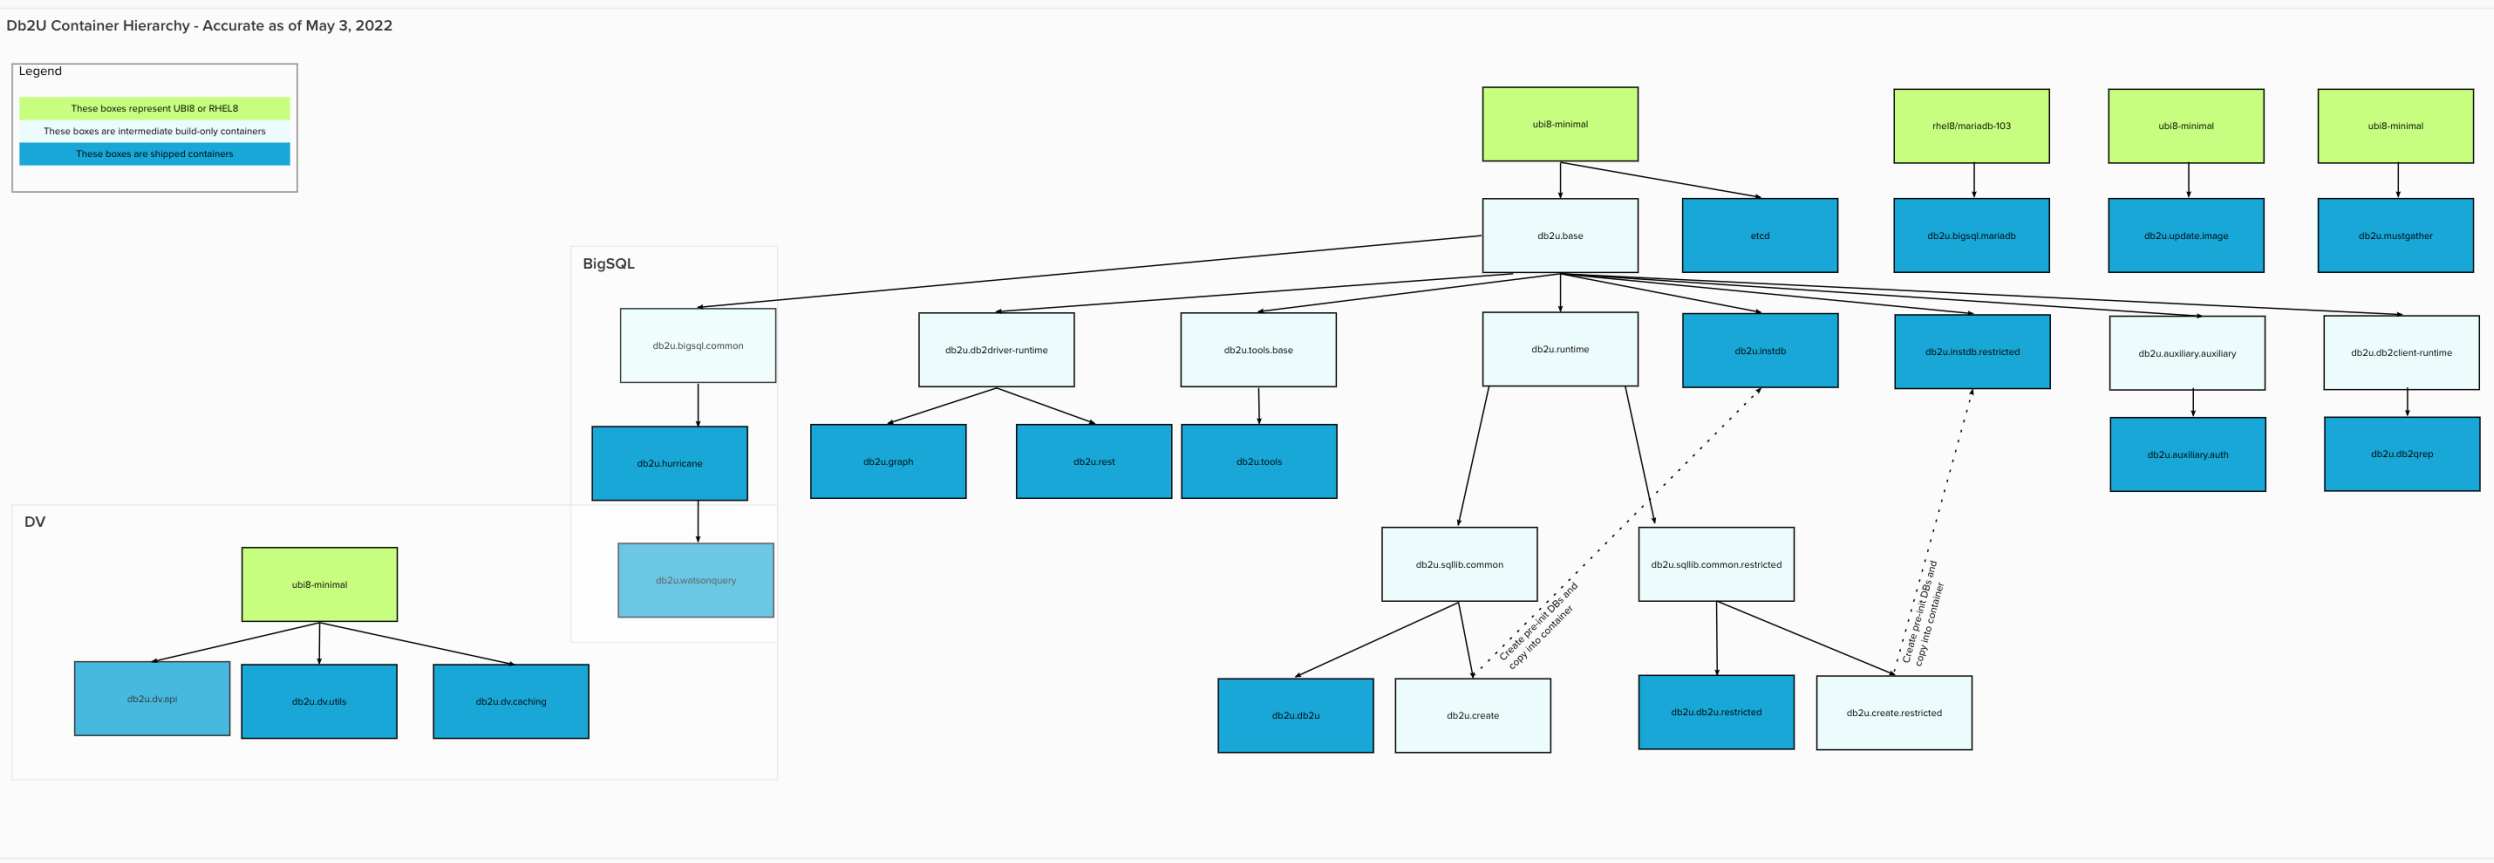

The Db2 “Go” Operator

- **Operational Management**
  - Package, Configure, Deploy & Manage
  - Deployment
  - Management
- Measured for completeness by a **maturity model**
  - Provides a glimpse at the Db2 Operator roadmap
    - Currently expanding Phase III capabilities
- **Delivery**
  - IBM Operator Catalog
    - Supports Air Gap
  - Red Hat Marketplace

The Operator **Maturity Model**

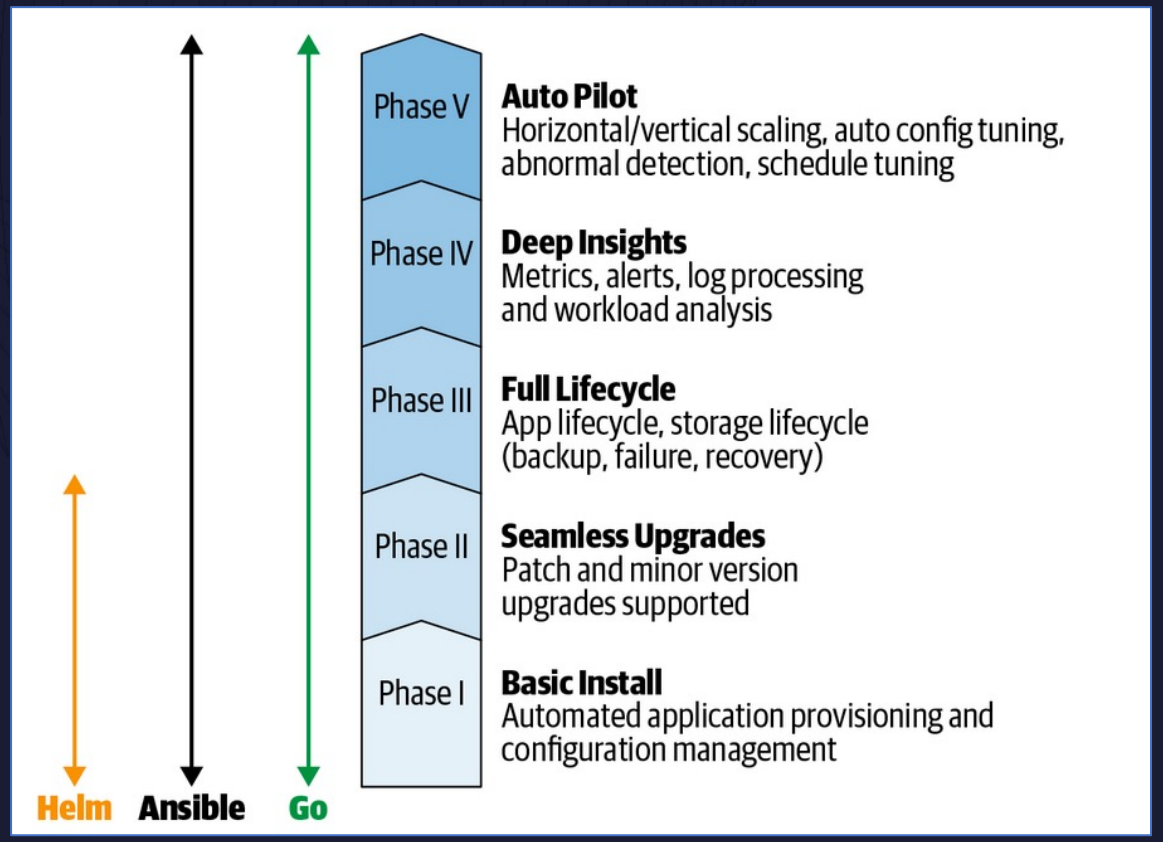

Db2U Architecture – Overview

- Underlying Kubernetes resource model:
  - Db2 Engine Pod lifecycle managed using a `StatefulSet` resource, since Db2 is a stateful application.
  - 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

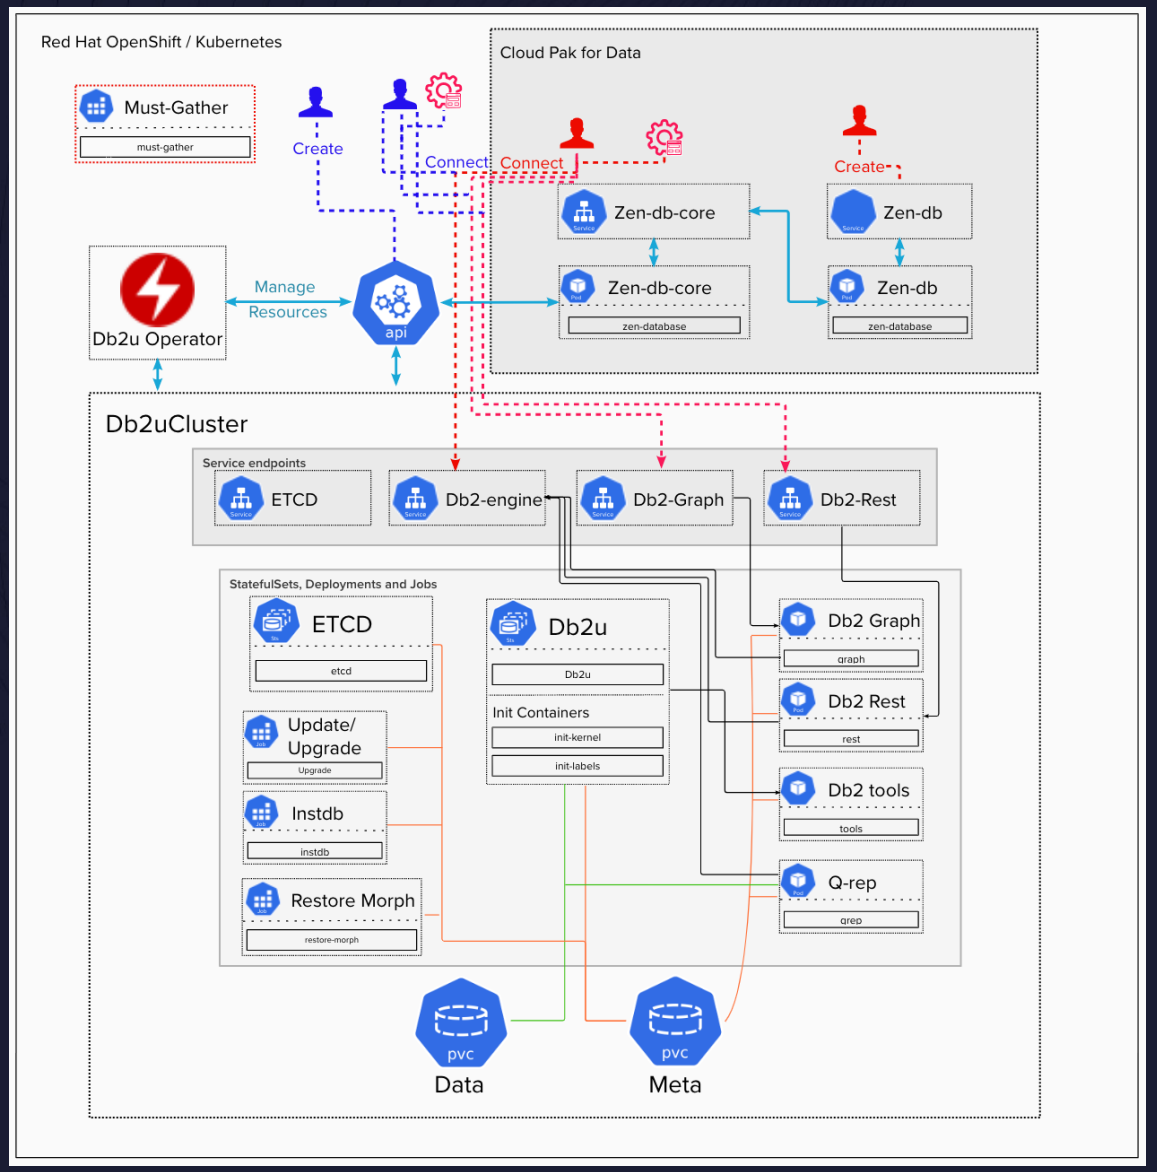

Db2U Architecture - Kubernetes Resource Model

- All Db2 configuration settings (Registry/DBM/DB cfg) injected via CR are transposed into `ConfigMaps` and mounted into Db2U PODs.
- Persistent Volume attachment:
  - Shared Storage volume (Db2 instance home/other shared metadata) via `PersistentVolumeCliaim` (PVC) with `ReadWriteMany` (RWX) access mode.
  - Data Storage (Db2 database paths) via `VolumeClaimTemplates` with `ReadWriteOnce` (RWO) access mode in Db2U `StatefulSet`

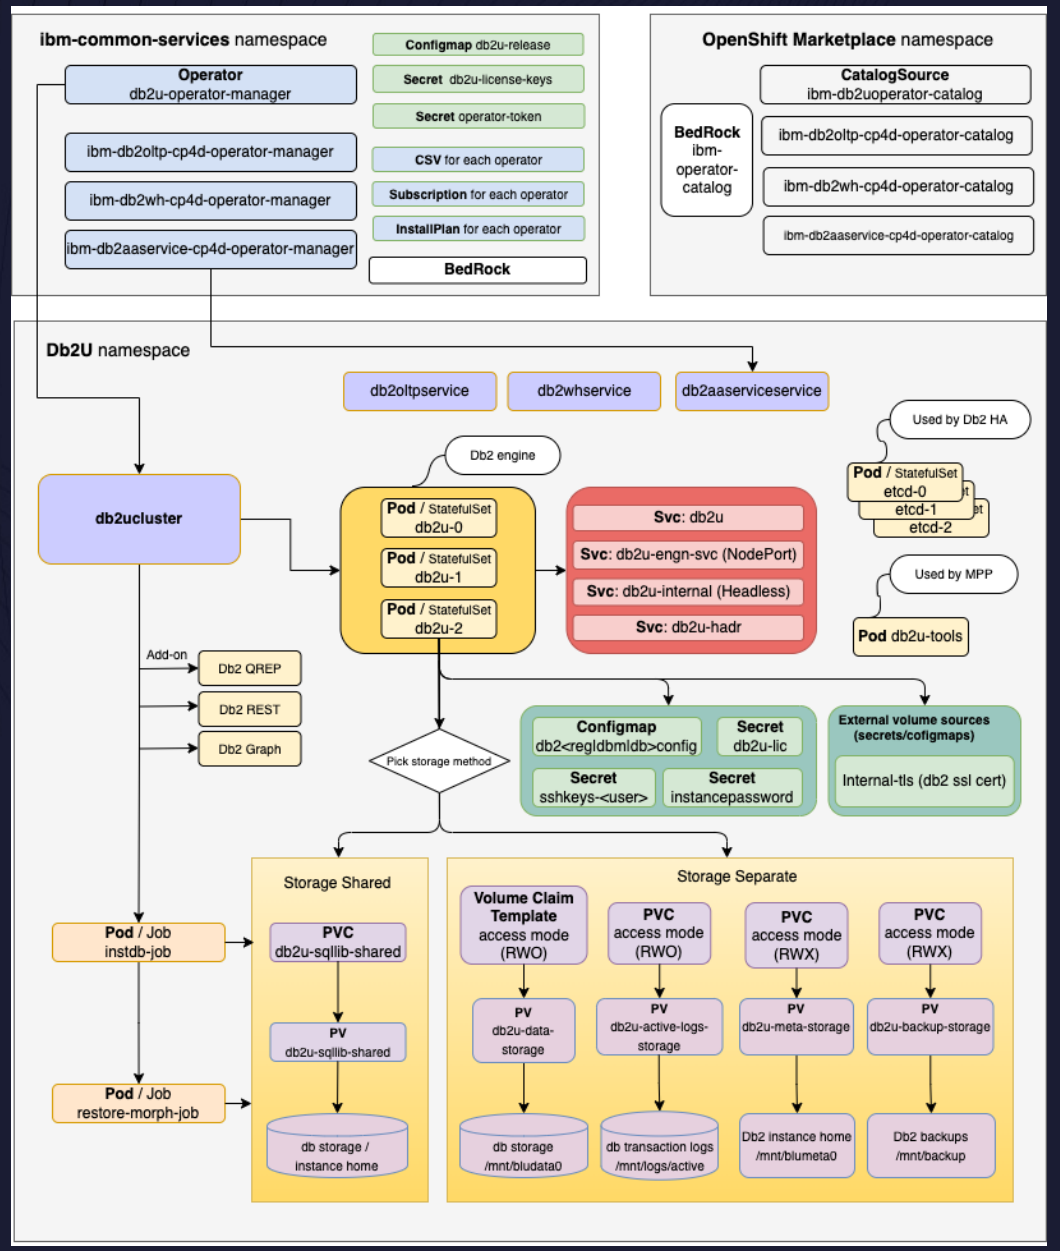

Db2U Architecture – Storage (1|3)

Cloud Native

- ❑ OpenShift Container Storage (OCS/ODF) 4.7+, 4.9
- ❑ Portworx 2.7+
- ❑ IBM Spectrum Scale CSI 2.0+
- ❑ Public Cloud Provider Native Storage (EKS/AKS)
- ❑ NAS (Dell EMC Isilon, NetApp Trident CSI)

Traditional

- ❑ NFS
- ❑ Host Path (IBM Cloud, Dell EMC Isilon, Local…)

Options for Loading

- ❑ NFS
- ❑ Remote Load or RSYNC to Scratch
- ❑ Db2 CP4D Interface provides a Remote Load capability(S3 etc)
- ❑ S3 External Tables

Db2U Architecture – Storage (2|3)

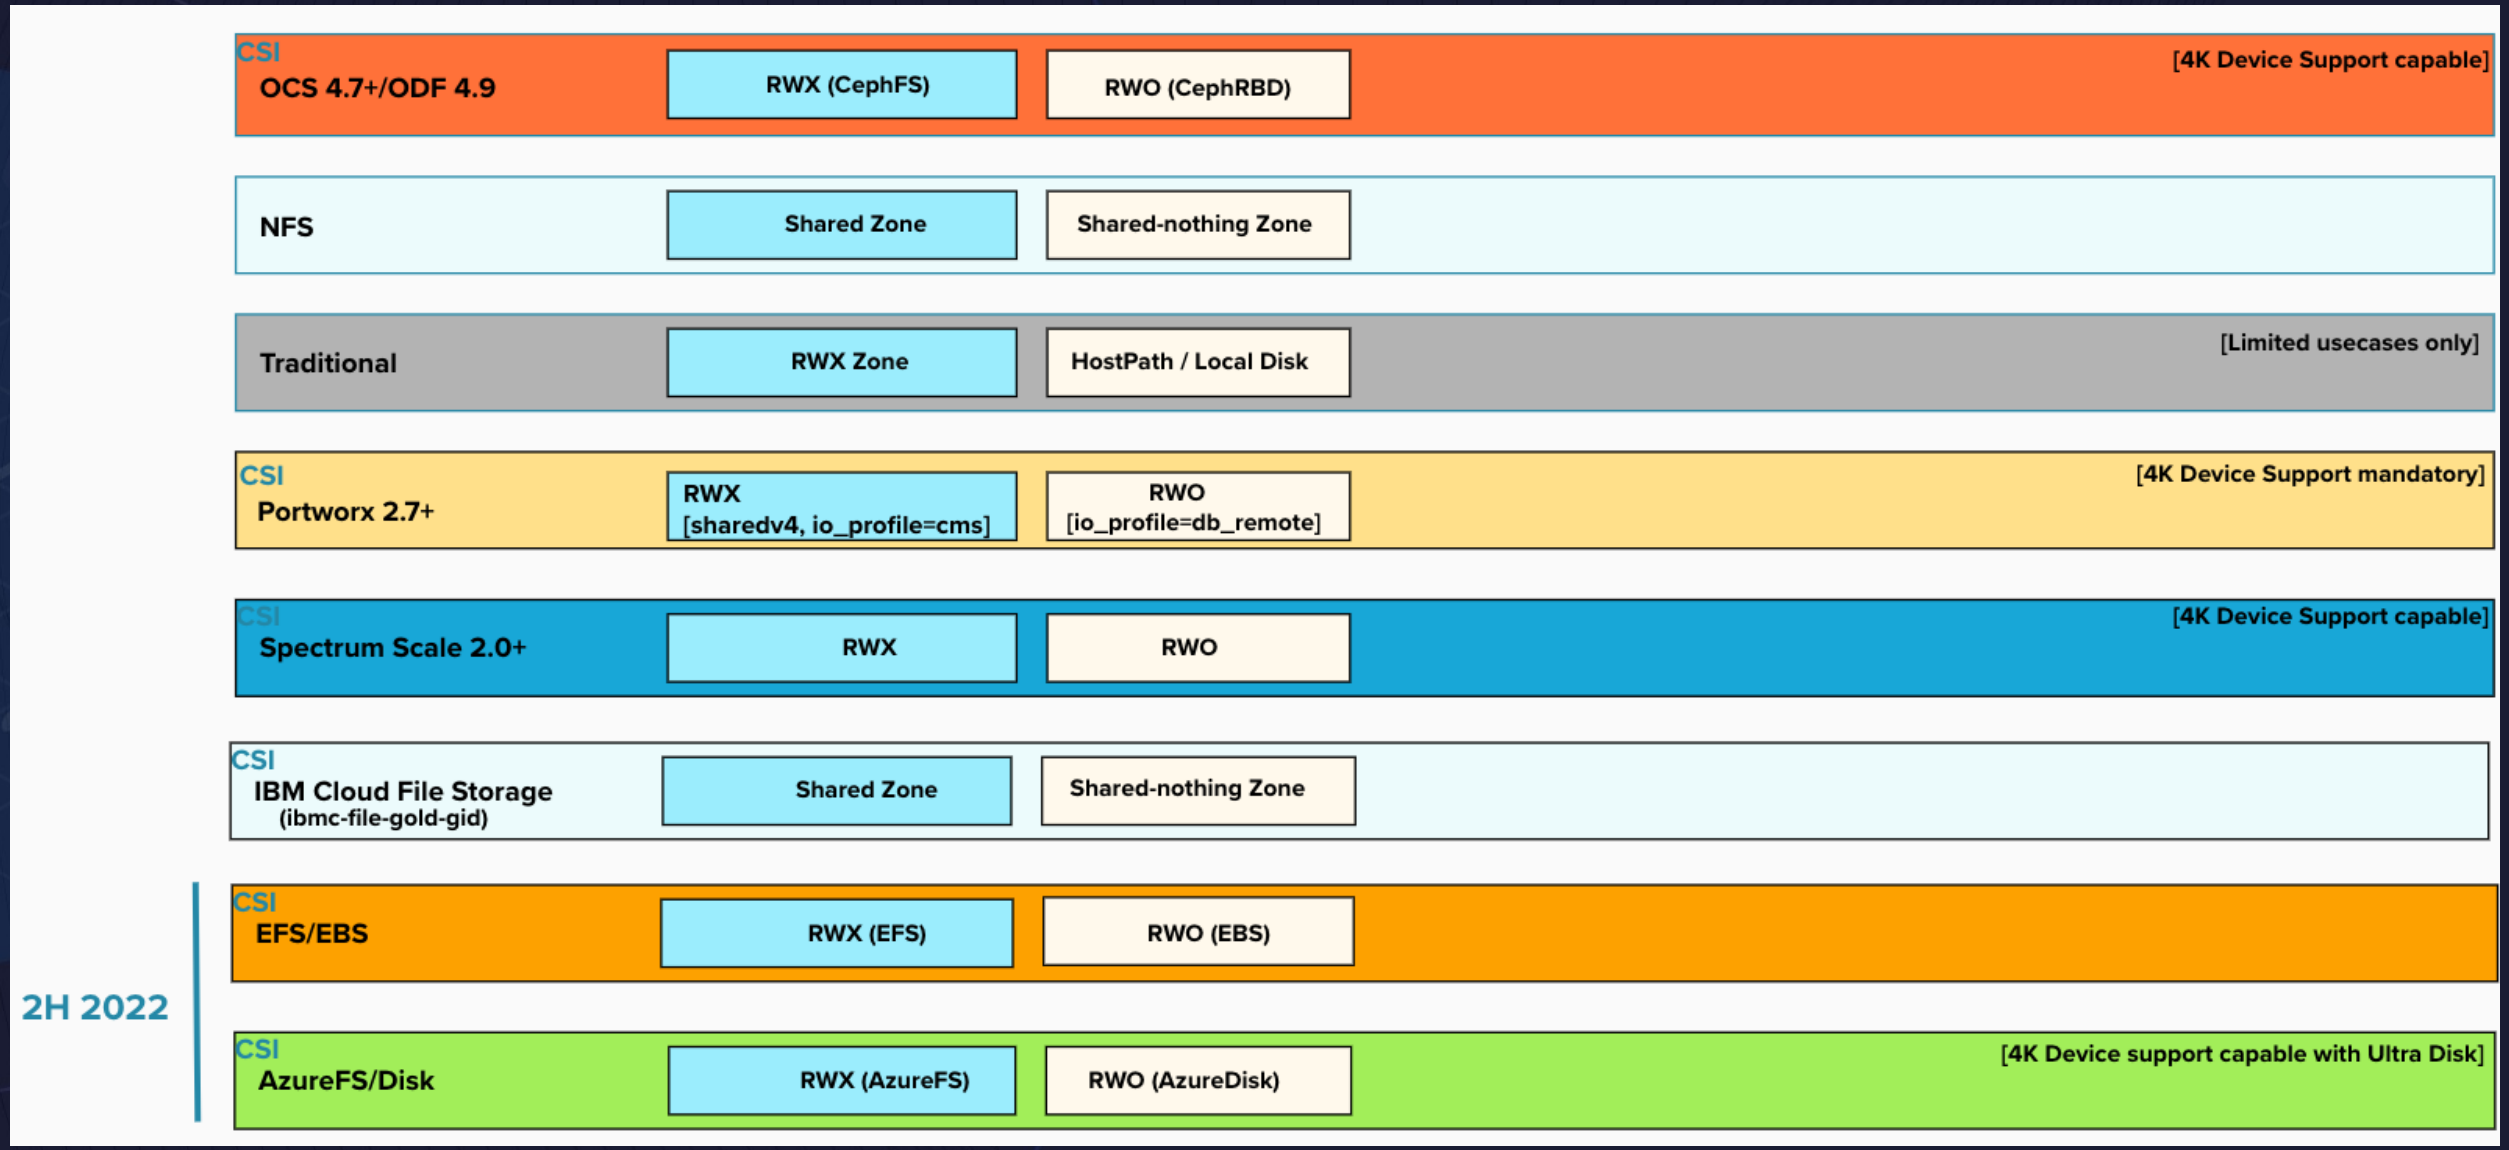

Db2U Architecture – Storage (3|3)

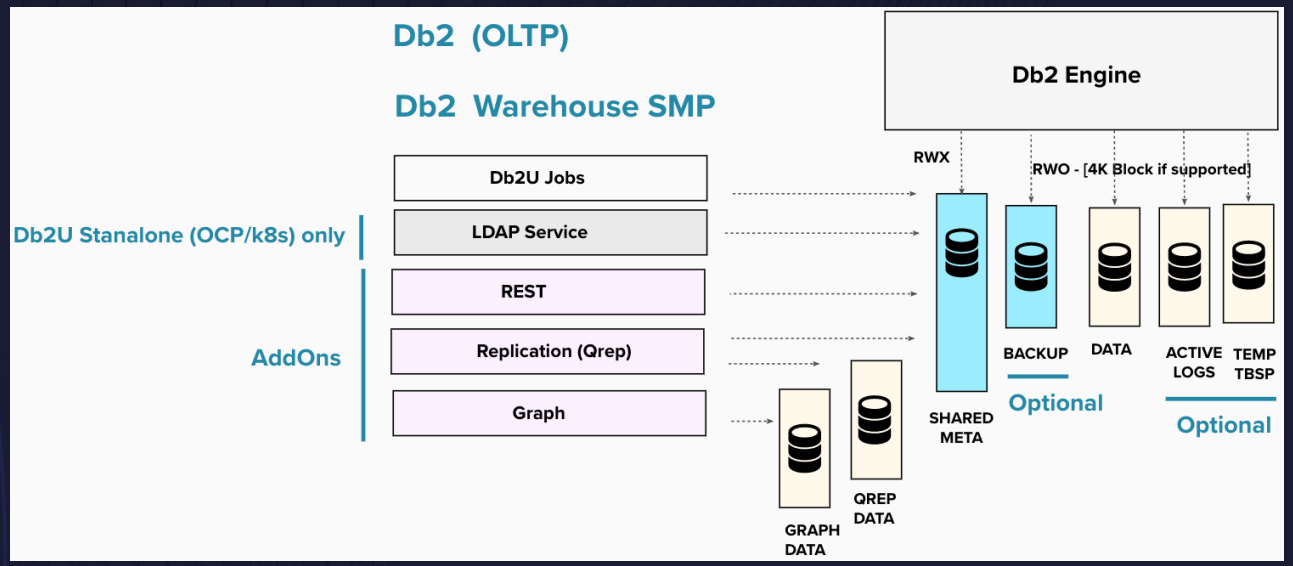

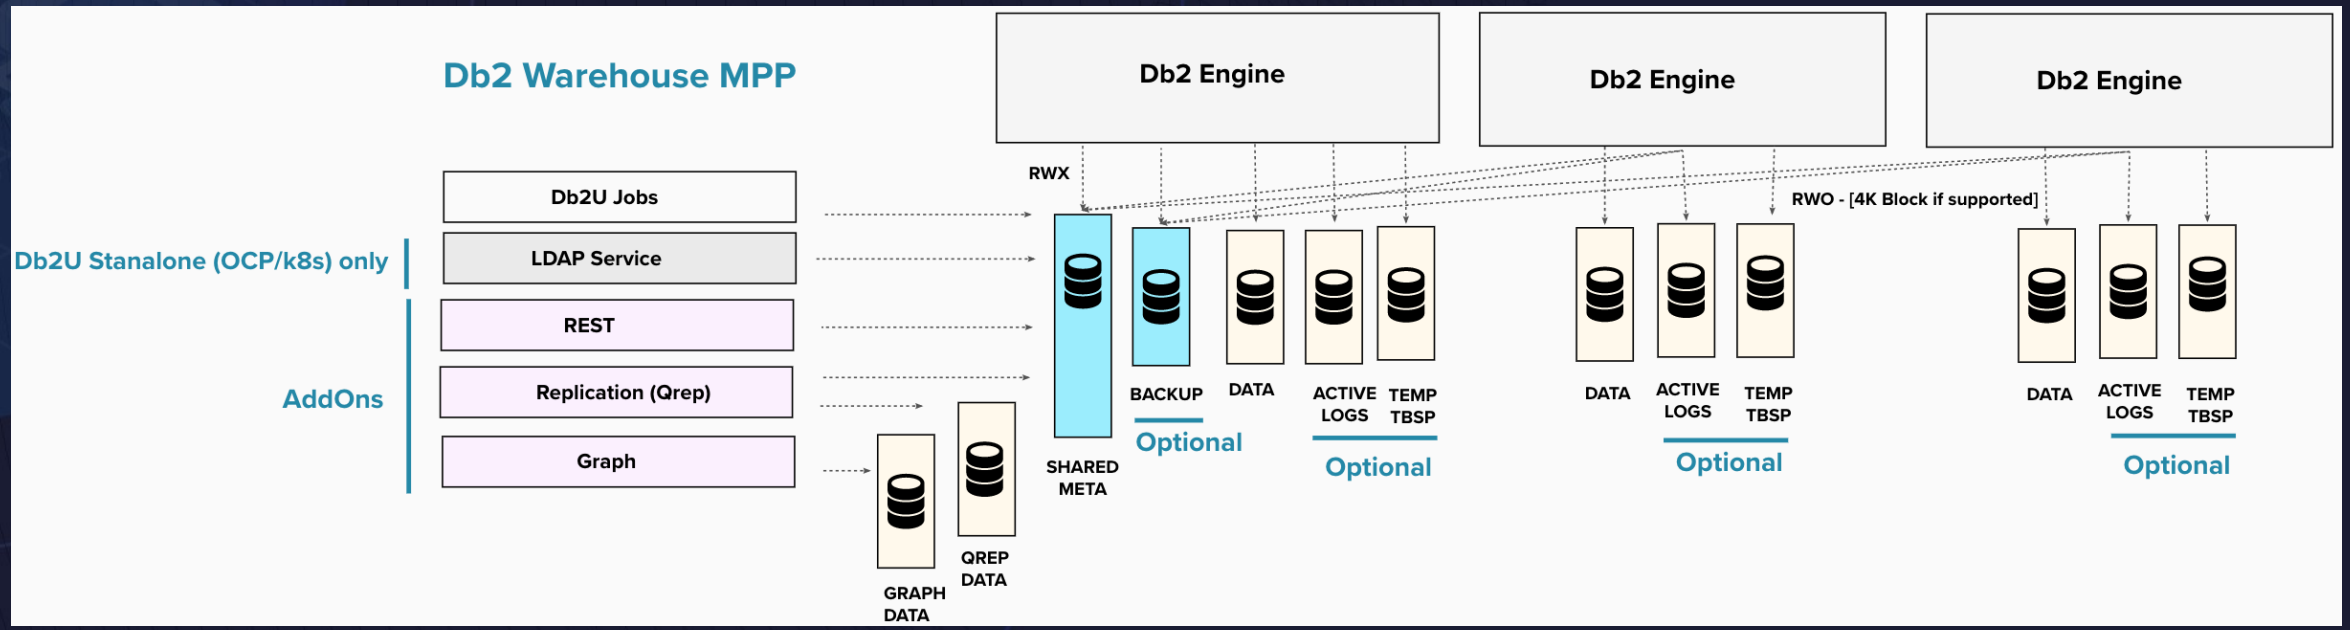

Db2 Universal Container (Db2U) Operators - Performance Evaluation

All Performance Results for **Db2U on RHOS**

Metrics	Db2U on OpenShift
Install Time	[Automated] ▪ 5’ – 10’ (Db2 OLTP / WH SMP) ▪ 10’ – 20 ‘ (WH MPP)
Upgrade Time	[Automated] 15’ – 30’ (Dependent on number of catalog objects)
[Db2 OLTP] HADR Takeover Time	[Automated] 2’ – 4’
[Db2 WH MPP] Recovery on Db2 Failure	[Automated] ▪ Non-catalog POD failure: 3’ + Db2 Crash Recovery time ▪ Catalog POD: 5’ + Db2 Crash Recovery time
Recovery on Hardware Failure	[Semi-automated] OCP/k8s Node failure detection time + *Recovery on Db2 Failure* time
Vertical Scaling Compute	[Semi-automated – will be fully automated soon] 2 -3’ + Db2 autoconfigure time

Reference architecture - https://www.redhat.com/en/resources/IBM-Db2-Warehouse-MPP-on-OpenShift-Container-Storage-detail

Db2 Operator - Value in all 4 quadrants

Modern
- HADR 4-Tier Resiliency
  - Self healing [Process & Pod]
  - Standby Node
  - Standby Db2 [HADR]
- Support for Container-Native storage i.e. OCS, Portworx, Spectrum Scale/Fusion CSI
- Resiliency w Replication
- Add-ons plugability

Managed Lifecycle
- Db2 Special Build
- Upgradeability - AutoUpgrade ["20']
- Kubernetes alignment

Infrastructure
- Better HW/VM utilization
- Deployment on AWS (EKS/ROSA) or Azure (AKS/ARO), aligned w DevOps
- API-driven management
- Automation through Operator
- Ready in ~5 minutes

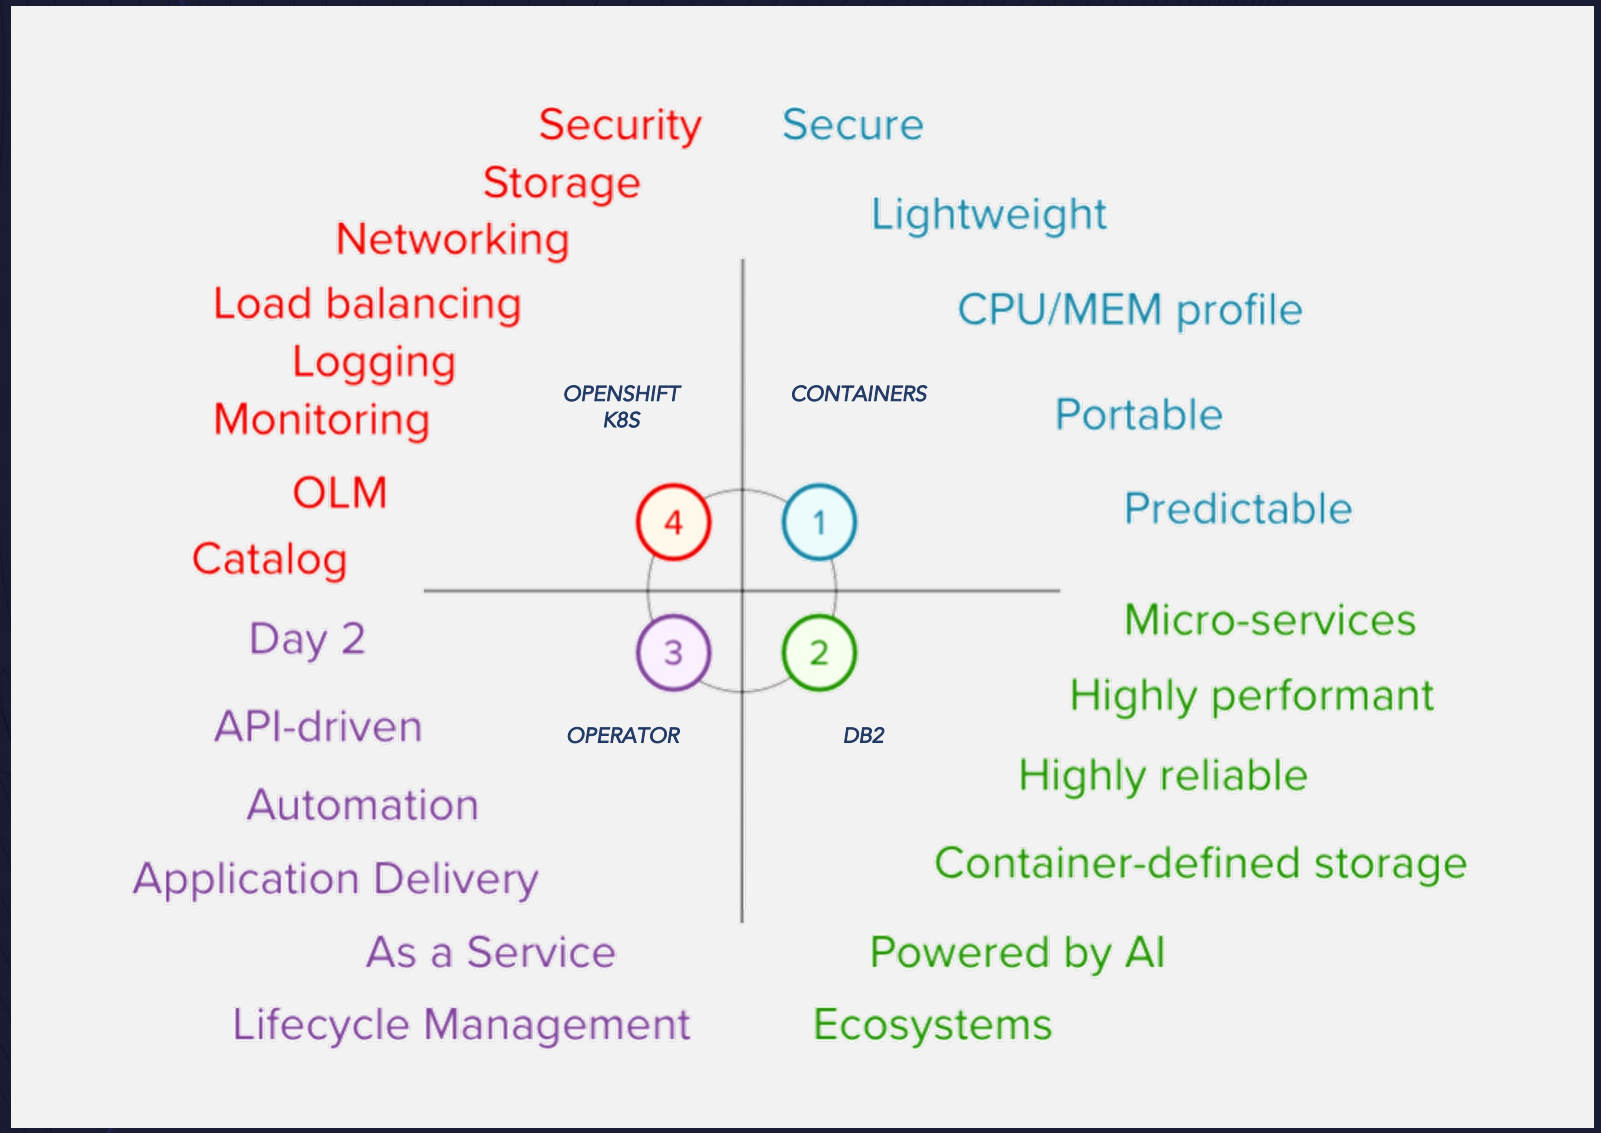

Db2U Next Generation: Core Capabilities (1|2)

- What's new
  - Db2 engine pods lifecycle is managed via a NEW Kubernetes Custom Resource - `Db2uEngine` Object
    - Spec of each Pod defined independently
  - One-to-One mapping between each database MLN storage path and Kubernetes volumes
    - Better alignment with MPP shared-nothing architecture
    - Leads to better support for horizontal scaling

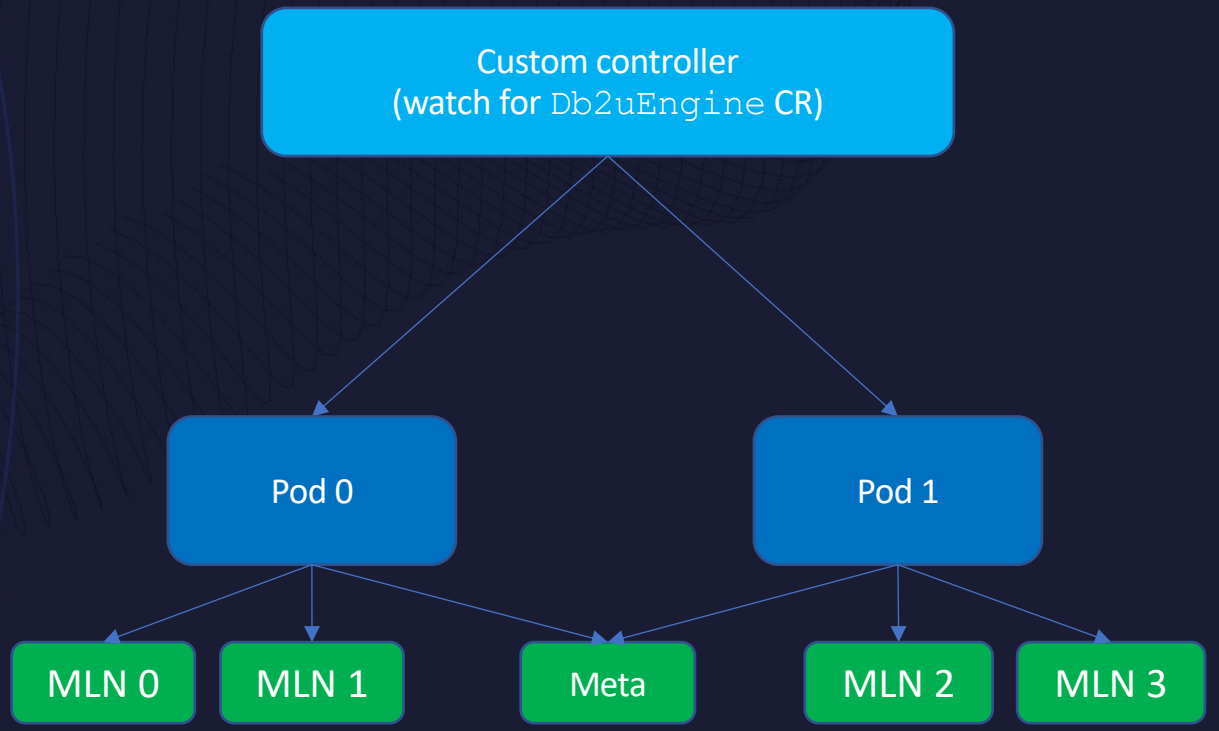

Db2U Next Generation: Core Capabilities (2|2)

- What's new
  - Ability to specify more than one volume for database storage for Db2 (OLTP) or Warehouse SMP deployments (1H 2023)
  - Grow database storage by adding a new volumes on-demand vs using CSI volume expansion (1H 2023)
    - Mitigate volume size limitations in cloud deployments (typically 16TB to 64TB cap depending on the vendor)
    - Better IO parallelism
    - Lower storage cost*

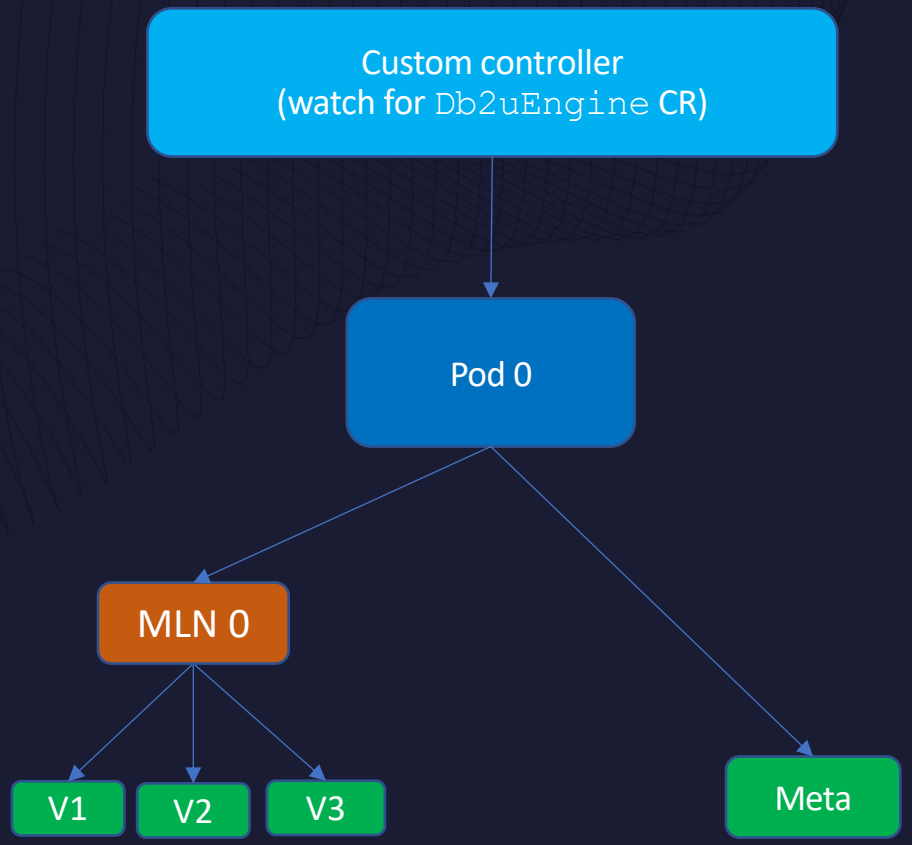

Next Generation: A Cloud-native Backup and Restore Experience (1|2)

[2H 2022] A Kubernetes controller driven approach to managing Db2 Backup/Restore, and Snapshot capabilities via Custom Resource Kind `Db2uBackup` and `Db2uRestore`

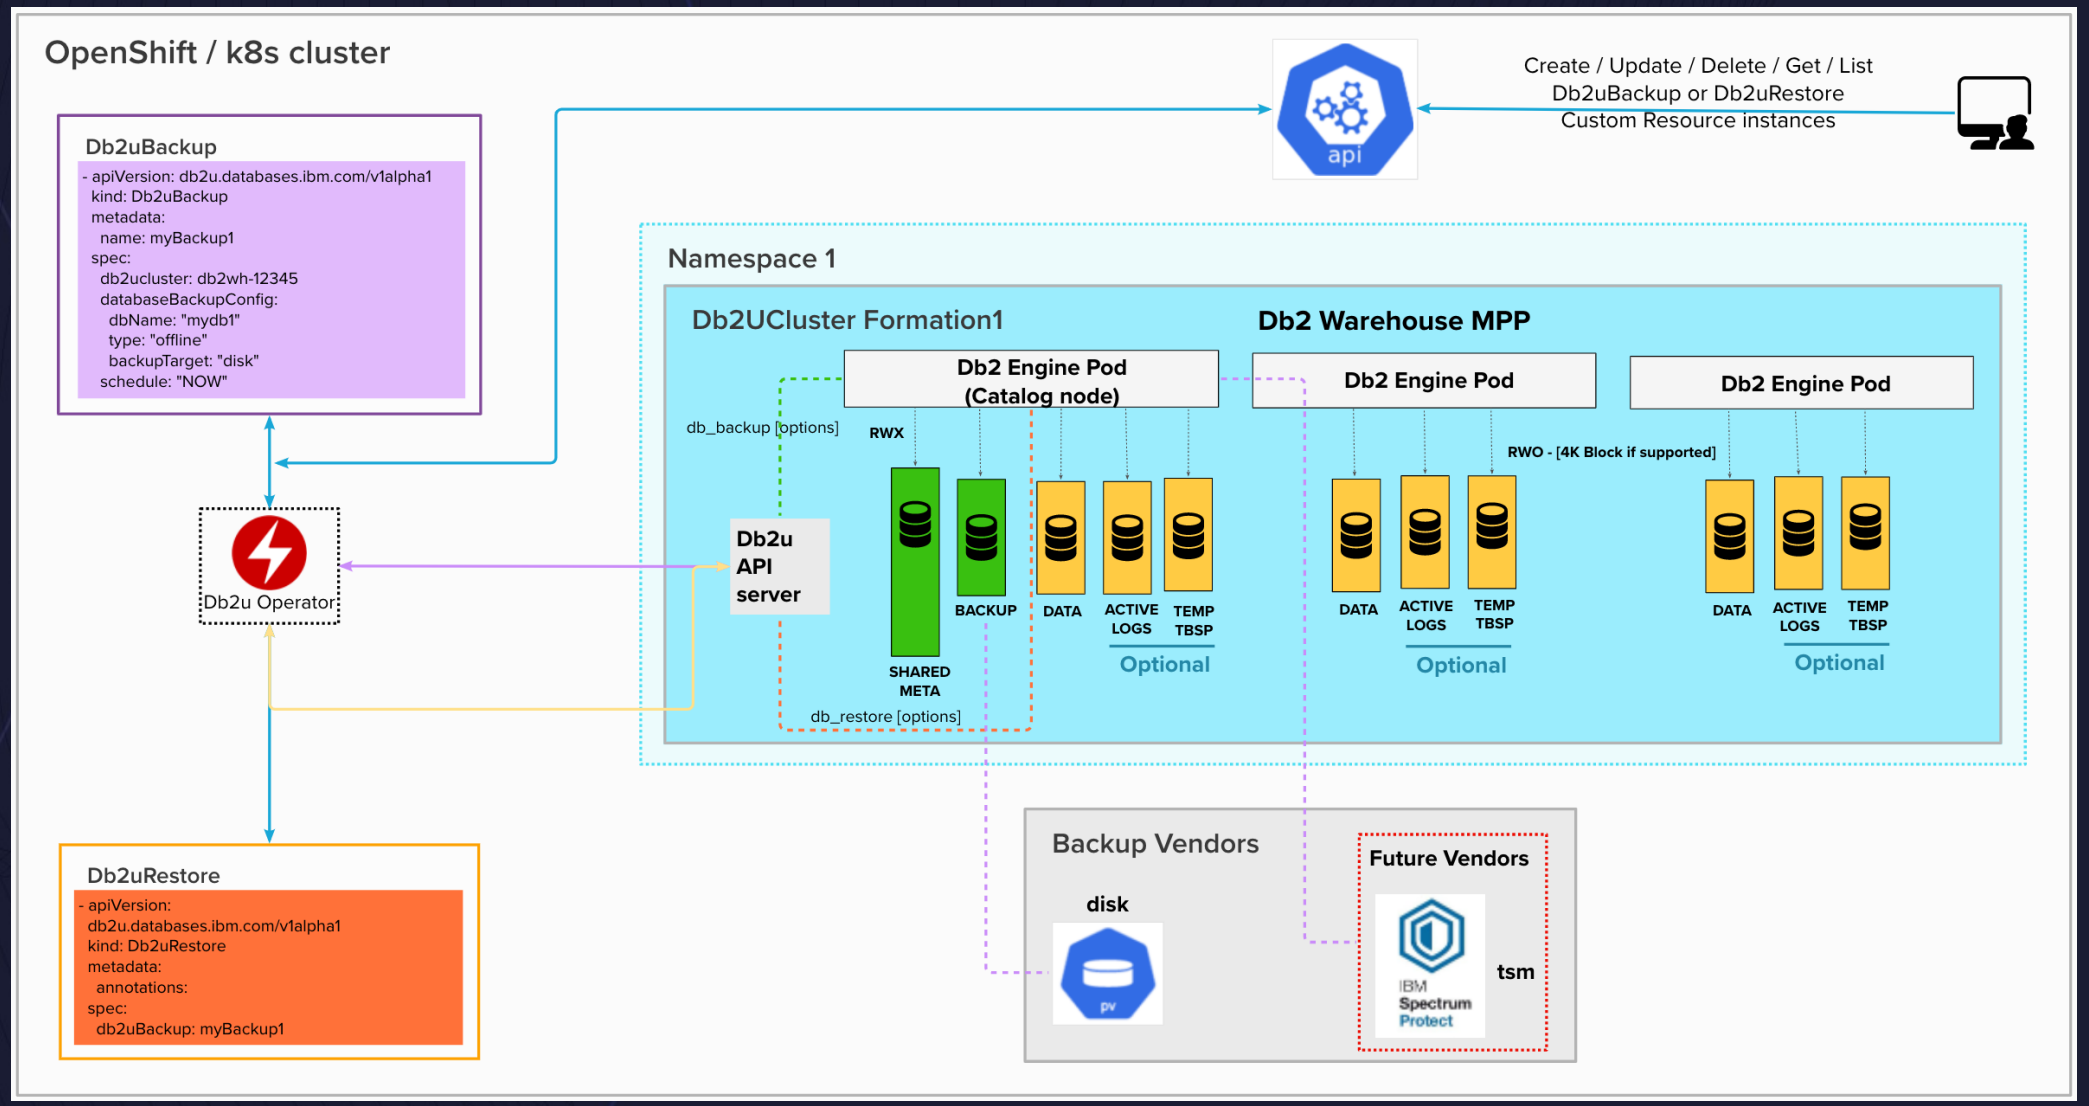

Next Generation: A Cloud-native Backup and Restore Experience (2|2)

Db2u Backup and Restore Custom Resource Definitions

```
- apiVersion: db2u.databases.ibm.com/v1alpha1
  kind: Db2uBackup
  metadata:
    name: myBackup1
  spec:
    db2ucluster: db2wh-12345
    databaseBackupConfig:
      dbName: "mydb1"
      type: "offline"
      backupTarget: "disk"
    schedule: "NOW"
```

```
- apiVersion:
  db2u.databases.ibm.com/v1alpha1
  kind: Db2uRestore
  metadata:
    annotations:
  spec:
    db2uBackup: myBackup1
```

```
- apiVersion:
  db2u.databases.ibm.com/v1alpha1
  kind: Db2uBackup
  metadata:
    annotations:
  spec:
    db2ucluster: db2wh-12345
    snapshotBackupConfig:
      dbname: "mydb1"
      excludeLogs: true
      volumeSnapshotClassName: "ocs-rbdplugin-snapclass"
    schedule: "NOW"
```

```
- apiVersion:
  db2u.databases.ibm.com/v1alpha1
  kind: Db2uBackupAndSnapshotSchedule
  metadata:
    name: myDb2BackupSchedule1
  spec:
    db2ucluster: db2wh-12345
    databaseBackupConfig:
      dbName: "mydb2"
      type: "online"
      backupTaget: "tsm"
    schedule: "0 12 * * *"
```

Next Generation: A Cloud-native Audit Facility

[2H 2022] Manage Db2 Audit facility via Db2U Custom Resource `addOns` mechanics:

- Apply the default audit policy
- Periodically archive audit records into audit tables in the database

Db2uCluster / Db2uInstance (next gen)

```
....
addOns:
  audit:
    interval: 15m
    applyDefaultPolicy: true
    archiveToDb: true
```

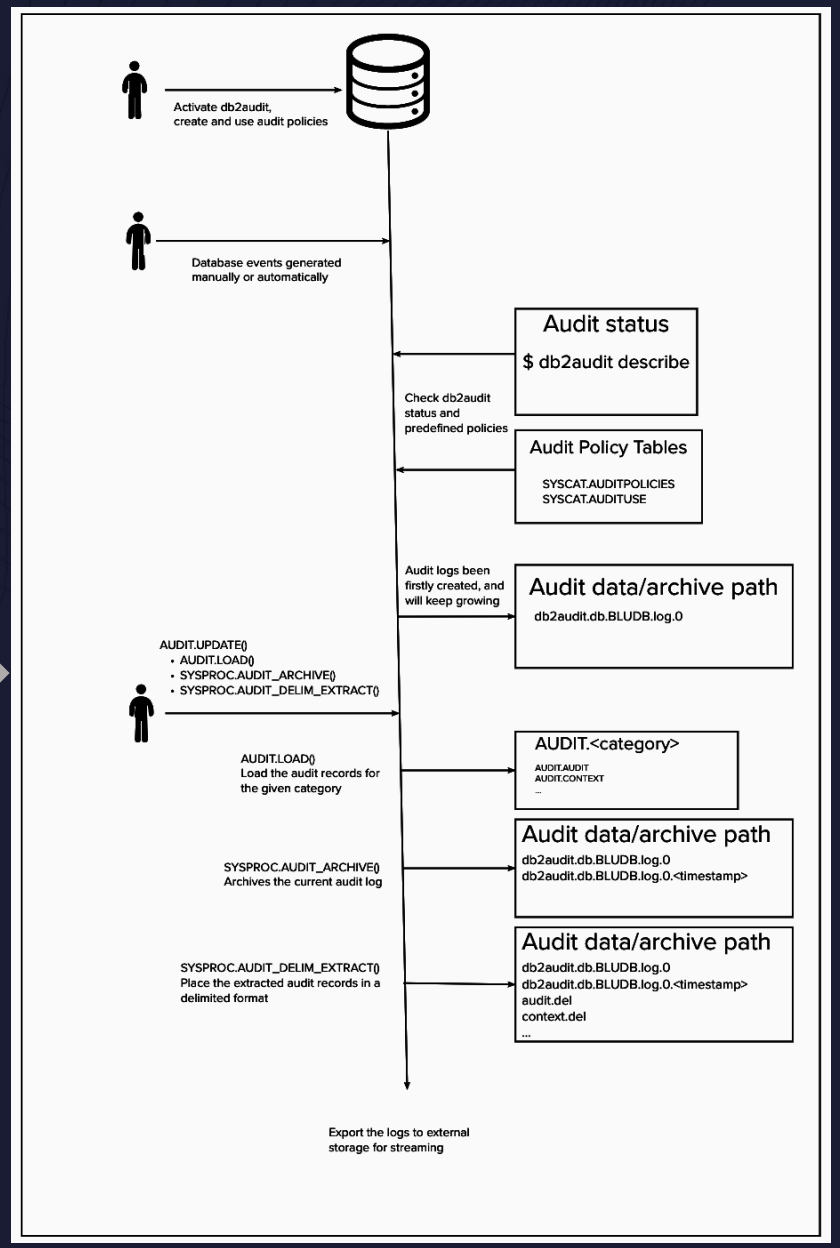

Next Generation: A Cloud-native Log Collection and Streaming

[1H 2023] Support log streaming (audit, diaglogs, HA logs, etc.) to Cloud Pak for Data or to an external logging service using a *Side Pod*, and enabled via Db2U Custom Resource `addOns` mechanics

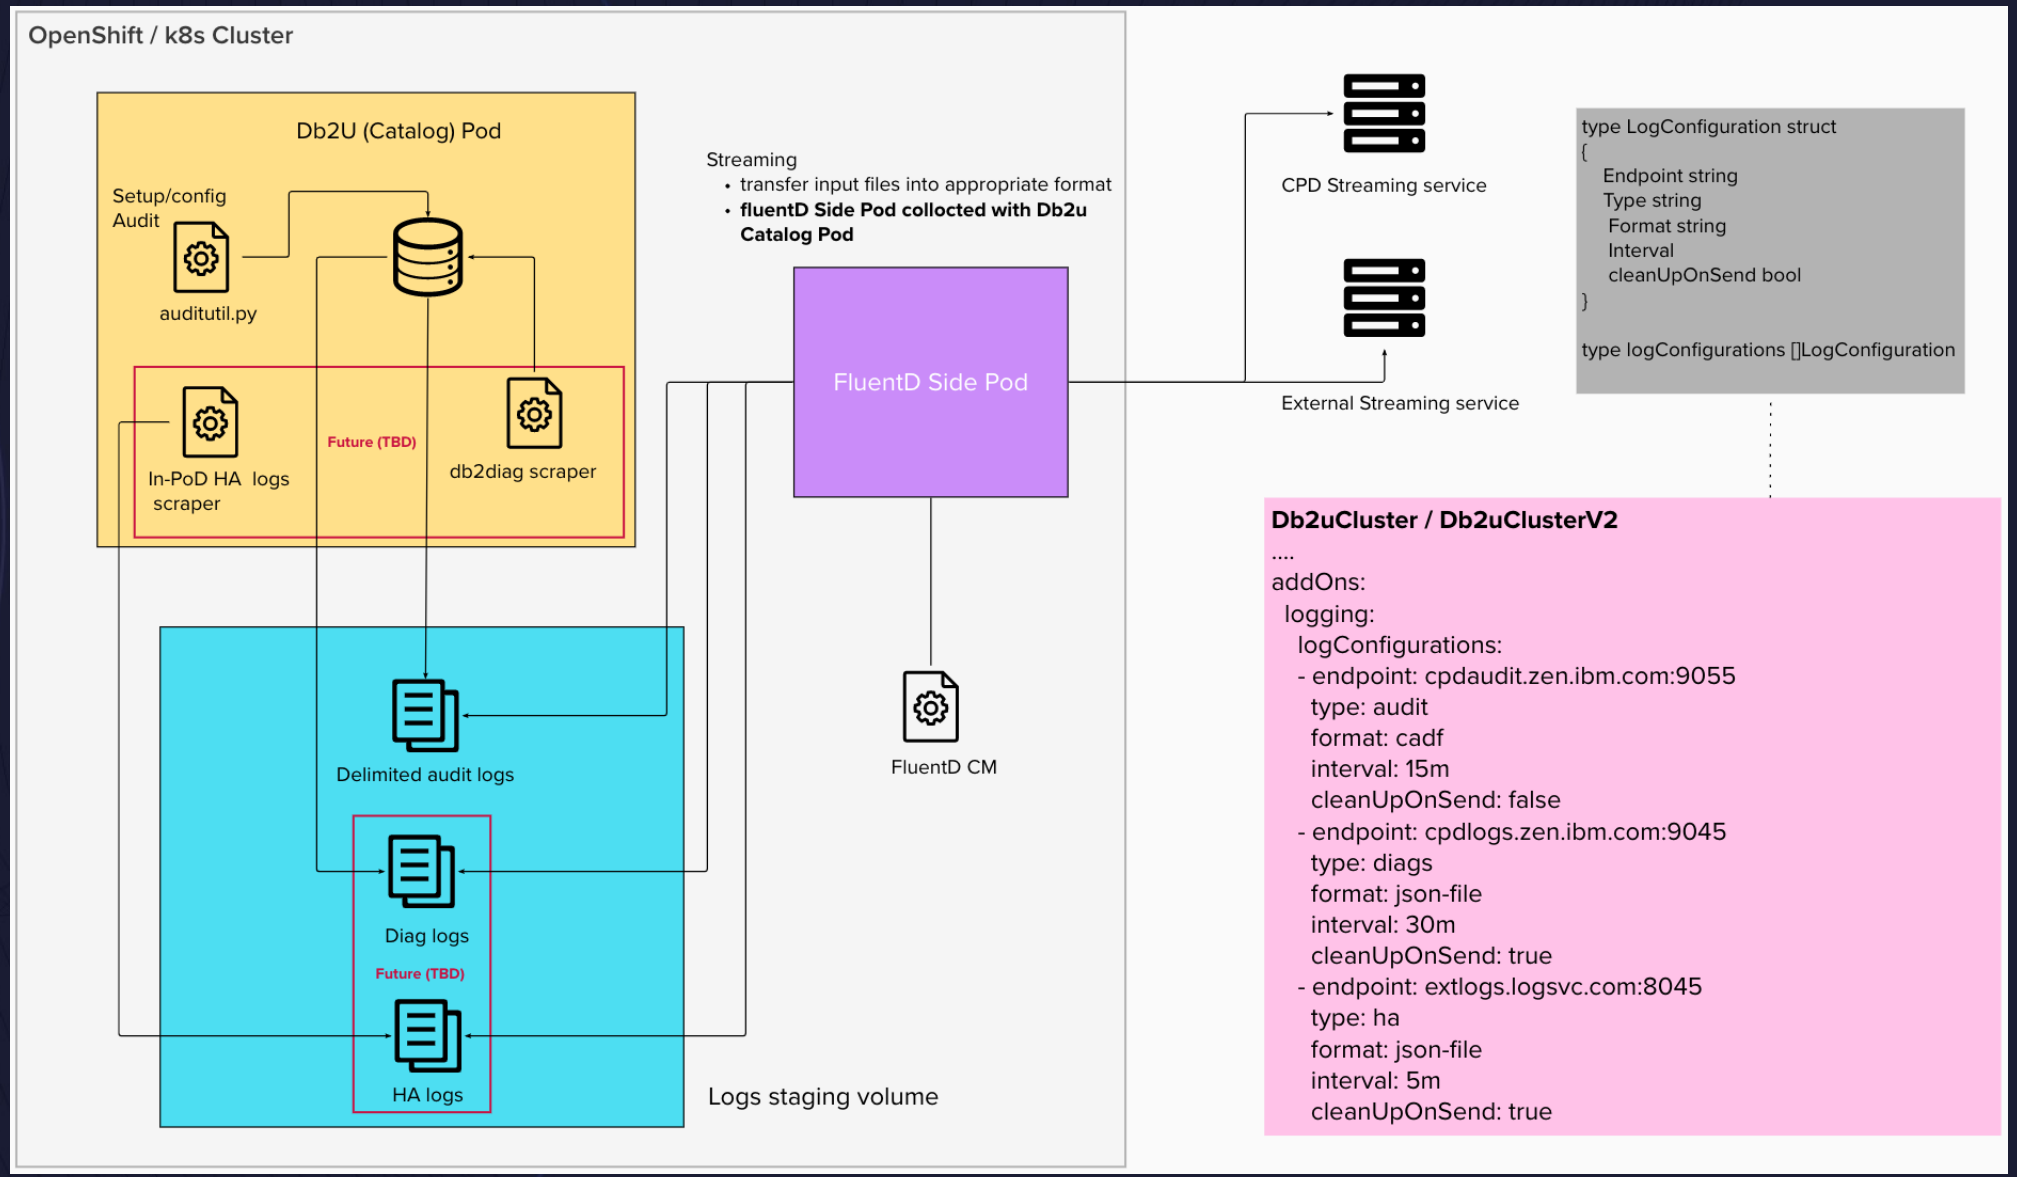

Db2U Next Generation – Perfectly Aligned for Public Cloud

- **Amazon**
  - Elastic Kubernetes Service (EKS) with EFS (Shared RWX) and EBS (Per-MLN RWO) volumes
  - Red Hat OpenShift Service on AWS (ROSA) with OCS/ODF
- **Azure**
  - Azure Kubernetes Service (AKS)
  - Azure Red Hat OpenShift (ARO)
- **Google Cloud** – Google Kubernetes Engine (GKE)

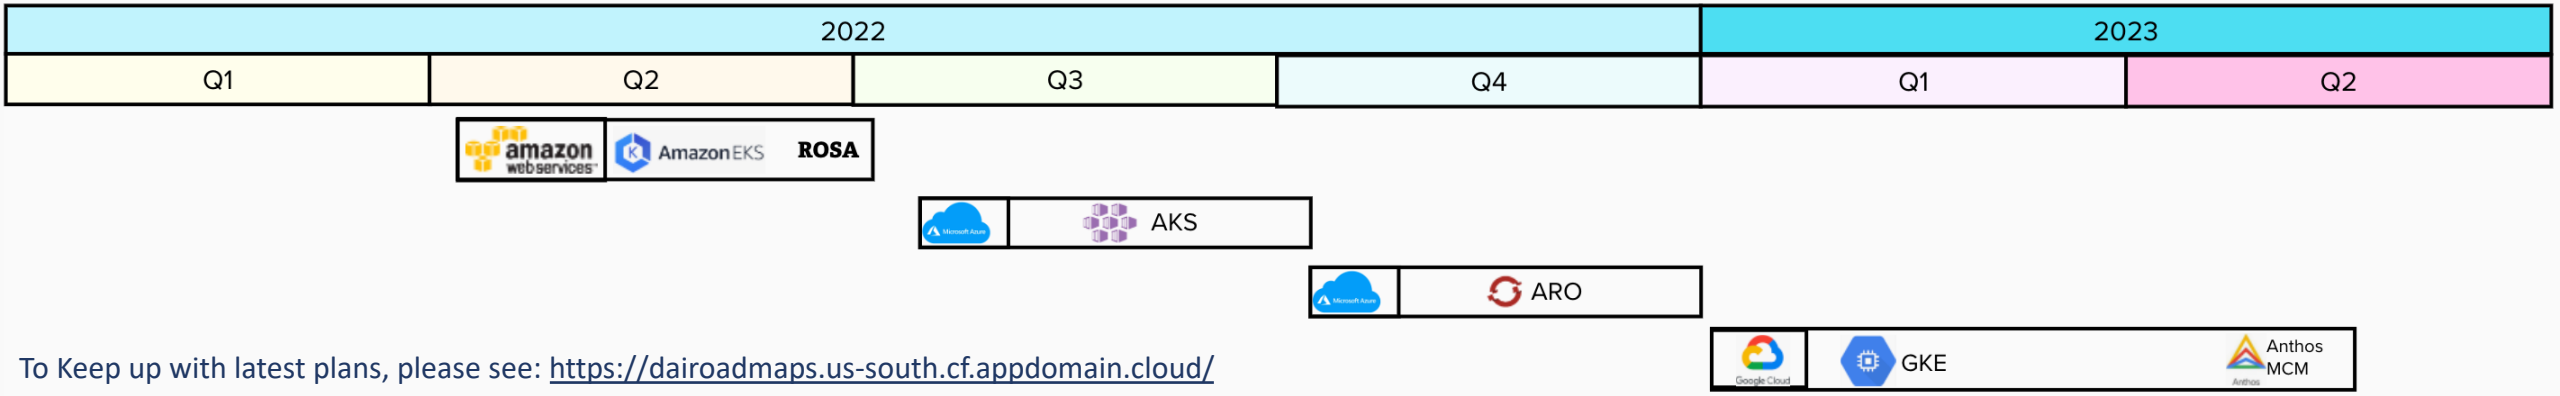

To Keep up with latest plans, please see: https://dairoadmaps.us-south.cf.appdomain.cloud/

Thank You

Speaker: Aruna De Silva, Architect – Db2U

Company: IBM

Email Address: adesilva@ca.ibm.com

Social:

- Linked-IN: https://www.linkedin.com/in/arunads/
- Twitter: @aruna_desilva