

# Kubernetes and Oracle: What Happened in 2024 and What's Next ?

Antonio Nappi CERN  
Miroslav Potocky CERN  
Adrian Karasinski



# The collaboration

Oracle collaboration has been ongoing, almost continuously, for nearly 25 years

## Collaboration evolved with time

- For many years focus on Oracle DB and WebLogic products
- Today it covers a wider scope
  - Anomaly Detection with ATLAS and CMS
  - Disaster Recovery on Cloud
  - Cloud Sustainability and FinOps
  - **Kubernetes**

# A bit of history

## What we do at CERN

- Provide hosting infrastructure to:
  - Java web applications (EDH, EDMS, impact, e-groups)
  - ORDS service (ADaMS, APEX)
  - CERN Single Sign-On service
  - WLCG Identity and Access Management service

## WebLogic shift to Kubernetes in 2020

- Started as R&D project in the context of Oracle openlab collaboration



**CERN Phonebook**



e-groups



# Recap from 2024 workshop

## From openlab workshop in 2024 [1]:

- Leverage statistical techniques to quantify the potential impact of a Kubernetes upgrade
  - Increase automatization
- Built a tool to spot API changes across Kubernetes(K8s) versions
  - Not all cases are covered
  - Unfortunately Kubernetes validation rules published as part of OpenAPI schemas may not be complete
- Follow up with CNCF and understand if there is something that can be improved

# Updates from 2024

## Meeting with CNCF Kubernetes Special Interest Groups (SIG)[2]:

- sig-api-machinery group

## Kubernetes developers are aware of the issue

- Agreed to increase awareness of the problem
- Pull request merged to official K8s docs

### OpenAPI V2

The Kubernetes API server serves an aggregated OpenAPI v2 spec via the `/openapi/v2` endpoint. You can request the response format using request headers as follows:

Header	Possible values	Notes
Accept-Encoding	gzip	<i>not supplying this header is also acceptable</i>
Accept	application/com.github.proto-openapi.spec.v2@v1.0+protobuf	<i>mainly for intra-cluster use</i>
	application/json	<i>default</i>
	*	<i>serves application/json</i>



### Warning:

The validation rules published as part of OpenAPI schemas may not be complete, and usually aren't. Additional validation occurs within the API server. If you want precise and complete verification, a `kubectl apply --dry-run=server` runs all the applicable validation (and also activates admission-time checks).

# Side projects

## Oracle Database Multilingual Engine (MLE)

- Evaluation of the MLE solution, encompassing performance metrics, stress testing, and analytical reporting
- Understand how MLE could fit and benefit the CERN environment

## ORDS - Oracle REST Data Services

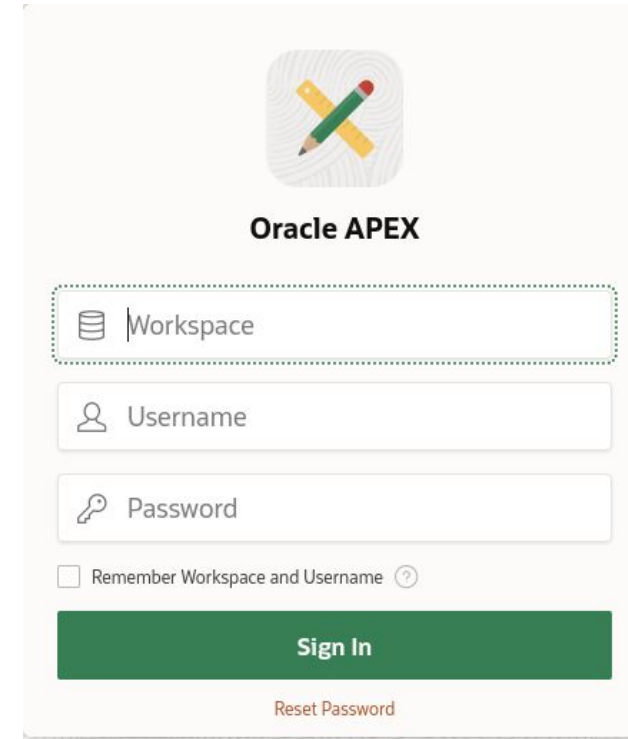
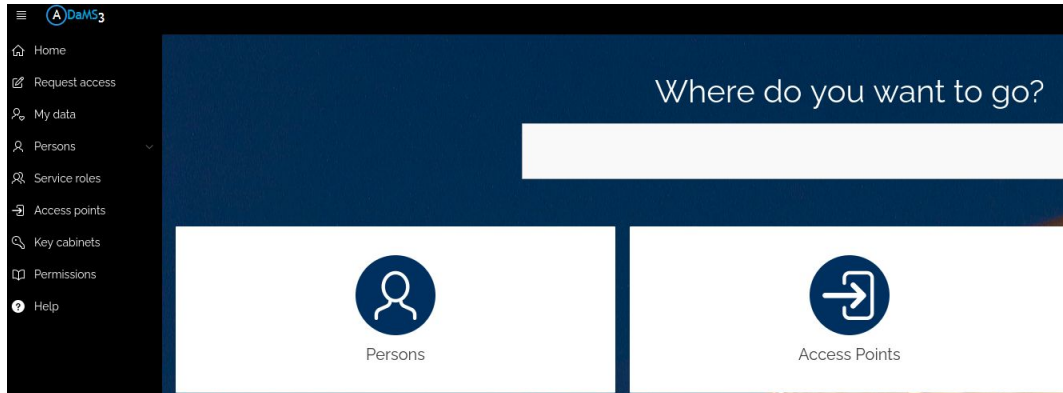
- Widely used at CERN, e.g. in Access Control (ADAMS)
- Refactoring, documenting and migration plans

## ORDS JWT (Json Web Token)

- Collaborating with Oracle ORDS developer team
- Testing role-to-privilege mapping using custom OIDC JWT claims

# Plan for Phase VIII

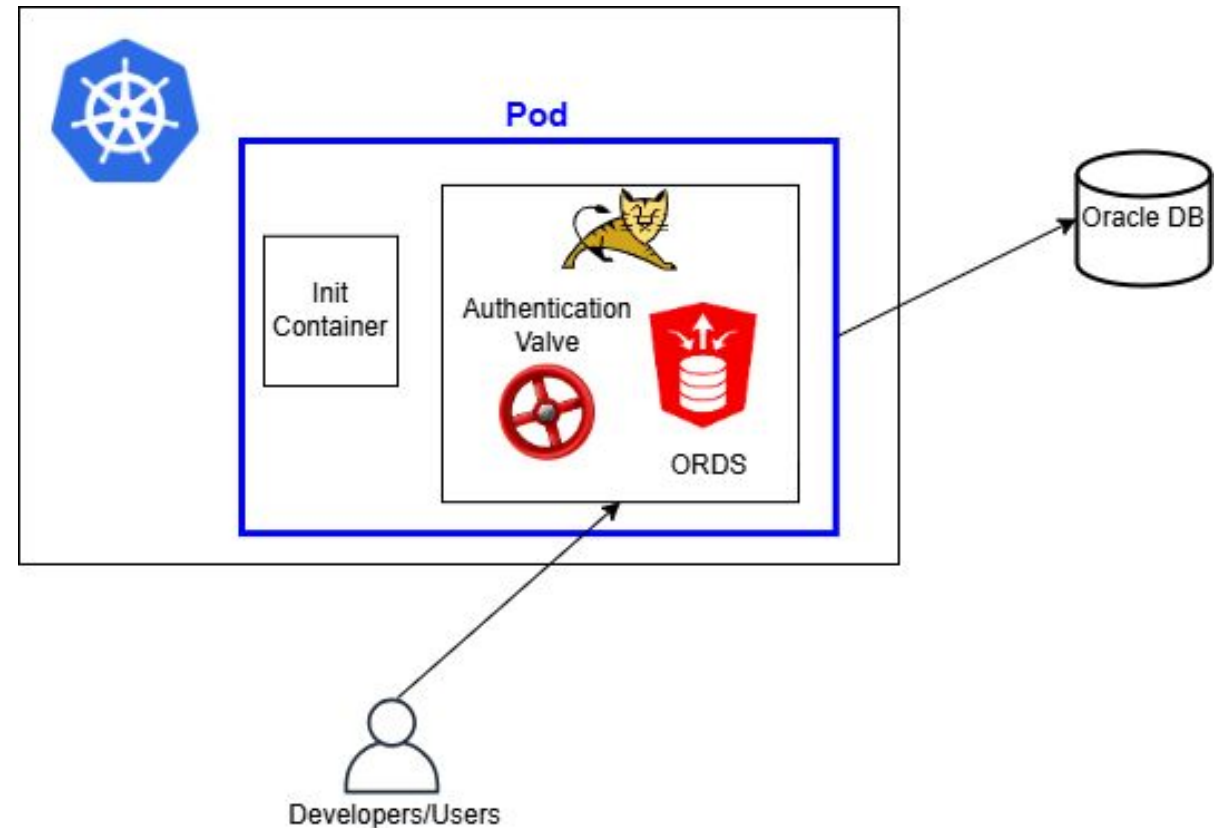
## Modernize Oracle REST Data Service (ORDS) and Oracle Database Service via Kubernetes



# ORDS service

## Tomcat Pod:

- Retrieves runtime configuration from an external database and builds a deployable WAR file.
- Provides custom authentication and authorization solutions.
- Integrates with CERN Single Sign-On (SSO).
- Translates CERN SSO group memberships into Java principal roles based on database definitions.





# ORDS service: Challenges

## Lack of standardization

- Custom authentication and authorization system to maintain and upgrade together with ORDS versions
  - Represents also a security risk
- Runs on K8s but not use standard approaches

# ORDS service: The Plan

## Replace custom deployment with ORDS Kubernetes Operator

- Follow best practices in Cloud Native Computing Foundation(CNCF) community
- Defines ORDS configuration as Kubernetes Custom Resource Definitions (CRD)

## Get rid of Tomcat assuming the ones below are bypassed:


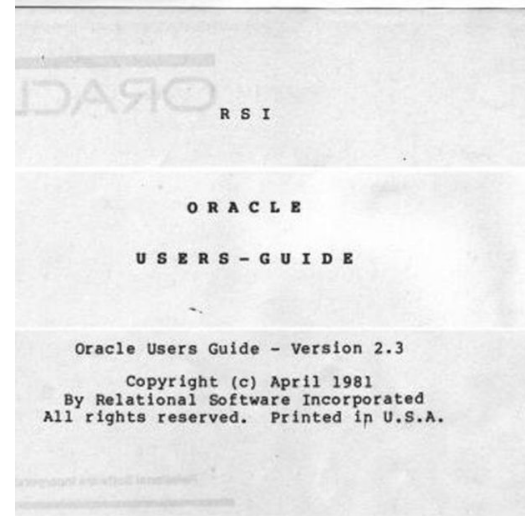
- Authentication and Authorization Valve
- Translation to CERN roles

## Leverage GitOps to trace and monitor changes

# Oracle Database service

## Oracle client since 1982

- 105 Oracle databases
- Almost 12,000 Oracle accounts
- Features snapshot
  - RAC, RAC One Node, RMAN, APEX
  - Active Data-Guard, ORDS, dNFS
  - KeyVault, Enterprise Manager



## Complex environment

- Five domains: Experiments, Accelerator, Administration, Engineering, General Purpose
- Hundreds of parallel database projects with dependencies

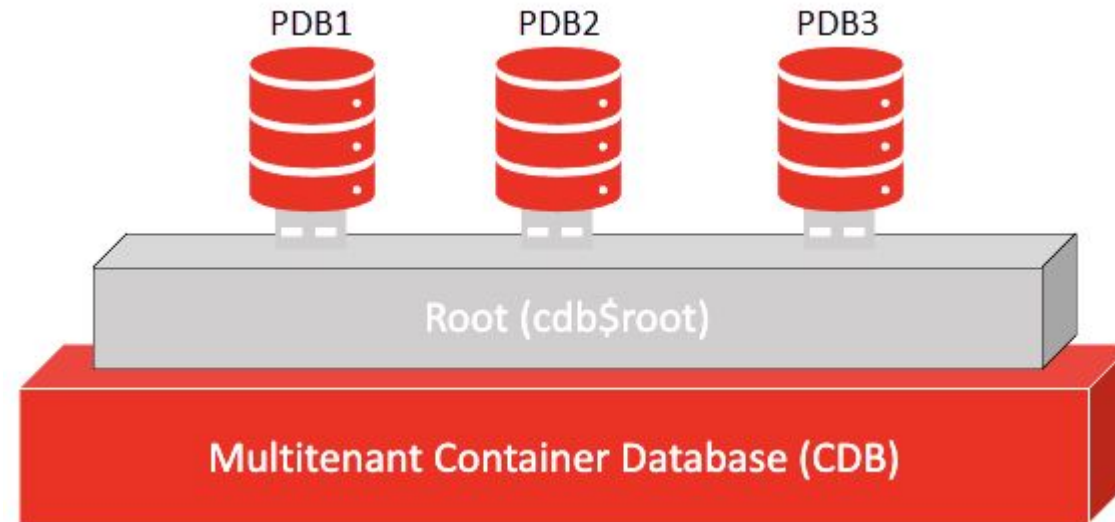
# Oracle Database service: Challenges

**From Oracle Database 21c multi tenant architecture becomes mandatory**

- CDB - container database
- PDBs - pluggable databases

**CERN Oracle Database service highly impacted**

**Oracle DB is becoming more integrated with the Kubernetes ecosystem**



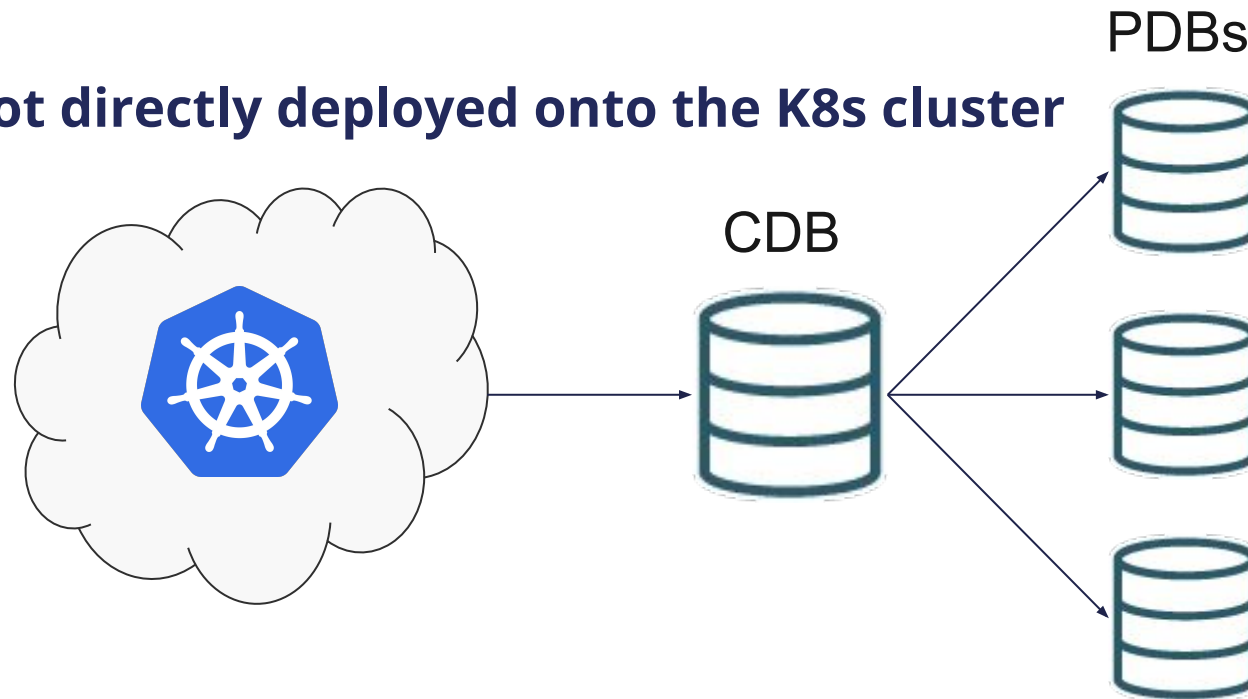
# Oracle Database service: The Plan

Empower developers to provision PDB instances within an existing CDB

## Management lifecycle of PDBs

- Creation, Deletion, Upgrade, Patching, Cloning, Migration

DB instances are not directly deployed onto the K8s cluster



# Modernizing with Kubernetes - Why

**Optimized Production:** Using Kubernetes best practices and Oracle Kubernetes Operator

**Improved Agility & Scalability:** Through Kubernetes and GitOps implementation

**Lower Operational Costs:** Achieved with automation

**Demonstrated Success:** Strong customer reference and community contribution

# References

[1] Migration between Kubernetes versions doesn't have to be error-prone?

[2] <https://www.youtube.com/watch?v=D0YWVZzv5-4>

# Thank you!

[antonio.nappi@cern.ch](mailto:antonio.nappi@cern.ch)

