



Oracle projects status and plans

CERN openlab Technical Workshop

Eva Dafonte Perez

12 January 2018

Overview

Database & Oracle Cloud
Oracle Management Cloud
Oracle Cloud Infrastructure
Analytics

Database & Oracle Cloud

New releases, new features

- Focuses on the deployment, evaluation through beta programs
- Non-volatile random-access memory (NVRAM)
- How far ingest and query workloads can be scaled for the upcoming LHC upgrade?
- Evaluate Oracle Zero Data Loss Recovery Appliance (ZDLRA) in an environment like CERN

Oracle Database Cloud Exadata Service

- Evaluate (cost/performance) advantages with CERN intensive applications
 - Using real production workloads, standby copies, ...
 - Use cases: SCADA system (control and monitor hw elements LHC and experiments) and QPSR (Quench Protection System) databases

Database & Oracle Cloud

Disaster Recovery solution in the Oracle Cloud

- Alternative for hosting standby databases (replacement of CERN remote data centre)
 - Data protection and availability
 - Offloading read-only workload and backups from production database
- Solutions for applications Oracle Container Cloud
- Functionalities and capabilities already tested in the previous phase, next is to assess the effort to integrate with the CERN environment and requirements:
 - Transferring data to cloud, integration with CERN network
 - Ability to switchover (planned event) or failover (unplanned event)
 - Ability to offload read-only production workloads to the cloud
 - Evaluate security constraints, integration with CERN SSO
 - Portability among cloud providers, fast and simpler provisioning
 - Monitoring

Oracle Management Cloud

Suite of integrated monitoring, analytics and management (on-premises and cloud) solutions.

Extend the tests to the different services provided

- Databases and applications deployed in Oracle Cloud
- Kubernetes clusters
- Streaming platforms like Apache Kafka
- Configuration management systems like Puppet

Explore and understand the monitoring and analytics features

- Replace our in-house monitoring tools
- Help to develop a Disaster Recovery plan for CERN applications
- Investigate how to anonymise data in the logs
- Identify areas/functionalities/features which could be improved

Oracle Cloud Infrastructure (OCI)

Evaluate new use cases provided by experiments

- CMS GPU study
 - Accelerate particle tracking tasks using an implementation on GPUs, based on a cellular automaton algorithm
 - Use generative adversarial networks to model the interaction of individual particles
- ATLAS HPC study
 - ATLAS G4 Simulation
 - ATLAS data reprocessing

Assess the effort needed to migrate different database applications from CERN Data Centre to Cloud

Analytics

Oracle Model Ops

- Currently under development
- Integrate within CERN heterogeneous and high data volume environment
- Evaluate features and identify areas to be improved

Enterprise Parquet

- Columnar storage format, commonly used by data processing frameworks
- Oracle is currently working to enhance its performance and functionalities
- Investigate its use for CERN Big Data projects

Analytics

Big Data SQL (BDSQL) and Kafka Integration

- Test DBSQL as a real time query layer for data in Kafka pipelines
- Evaluate DBSQL as a transparent SQL based data access layer
- Deployment in CERN environment

PGX/Graph Analytics

- Assess data analytics advantages of graph databases over traditional systems
- Evaluate potential application of PGX technology within CERN context
 - Zenodo – Analysis of scholarly publication data
 - CERN Fault Tracking System: connections analysis system for CERN's accelerator faults

Analytics

Oracle Analytics Cloud (OAC)

- New PaaS (Platform as a Service) providing a complete and elastic Business Intelligence platform in the cloud, customizable and managed by the user
- Migrate current Oracle Big Data Discovery (BDD) developments to OAC
 - e.g. Future Circular Collider Study
- Explore new use cases



Questions?

eva.dafonte.perez@cern.ch