



Contribution ID: 42

Type: **not specified**

EOS as storage back-end for JRC scientific data processing

Monday, 4 February 2019 16:20 (20 minutes)

The Joint Research Centre (JRC) of the European Commission has set up the JRC Earth Observation Data and Processing Platform (JEODPP) as an infrastructure to enable the JRC projects to process and analyze big geospatial data in support to EU policy needs. This platform is built upon commodity hardware and has become fully operational mid 2016, with platform extensions in 2017 and 2018. EOS was deployed on the JEODPP as the core storage system and is used in production with strong support by the CERN EOS team.

The JEODPP infrastructure is actively used by more than 25 JRC projects for data storage and various types of data processing and analysis. In order to serve the growing needs for data storage and processing capacity by the JRC projects, the platform has been extended in 2018. It currently consists of the EOS system as storage back-end with a total gross capacity of 13 PB, and processing and service nodes with a total of 1400 cores. A further extension to 17 PB gross storage and 1800 CPU cores are already in the pipeline.

In 2018 the EOS service has been migrated to the new Citrine release. The presentation will give an overview about the implemented platform, the current status, issues identified during the migration to the Citrine release, and experiences made with EOS and FUSE client as main storage back-end. Testing of new functionalities, like the MGM namespace in QuarkDB, wider usage of the fusex client, or the workflow engine are envisaged for 2019.

Primary author: Mr BURGER, Armin (European Commission - Joint Research Centre)

Co-authors: Mr TOGNOLI, Pier Valerio (European Commission - Joint Research Centre); Mr EYRAUD, Franck (Contractor of European Commission)

Presenter: Mr BURGER, Armin (European Commission - Joint Research Centre)

Session Classification: EOS Ecosystems