



Contribution ID: 10

Type: **Lightning Talk (8' + 2' for questions)**

Batch Processing on Kubernetes using the OSCAR framework

Monday, 28 October 2024 17:30 (20 minutes)

OSCAR is an open-source framework built on Kubernetes (K8s) for event-driven data processing of serverless applications packaged as Docker containers. The execution of these applications can be triggered both by detecting events from object-storage systems, such as MinIO or dCache (asynchronous calls) or by directly invoking them (synchronous calls). OSCAR is being used in several research projects, for example, it is the main inference framework used in the AI4EOSC (Artificial Intelligence for the European Open Science Cloud) platform. Related with this project, iMagine, Imaging data and services for aquatic science, is using the AI4EOSC platform to support their eight marine use cases. In the context of the iMagine project, the need to process big amounts of files by AI (Artificial Intelligence) models on top of the project's serverless platform has arisen.

For this purpose, we have designed OSCAR-Batch [<https://github.com/grycap/oscar-batch>]. OSCAR Batch includes a coordinator component developed in Python where the user provides a MinIO bucket containing files for processing. This component calculates the optimal number of parallel service invocations that can be accommodated within the OSCAR cluster and distributes the image processing workload accordingly among the services. OSCAR Batch uses a strategy based on sidecar containers, where the content of the MinIO bucket is mounted as a volume and accessible for the instance of the service (a K8s pod), preventing from moving large amounts of files. With this strategy, OSCAR-Batch ensures the efficient use of available CPU and memory resources in the OSCAR cluster and facilitates accessing to the output results, that will be stored also in MinIO.

This tool is mainly intended to process many files, such as historical data from an observatory service. For example, coming from the iMagine project, we support the analysis of historical data (hundreds of thousands of images) coming from the OBSEA underwater observatory by means of an AI-based fish detection and classification algorithm based on YOLOv8. With OSCAR-Batch, use case owners can reanalyze their data each time a new version of the AI model is produced.

This contribution presents OSCAR-Batch and exemplifies its usage with real use case scenarios coming from the iMagine project.

This work is partially funded by the project iMagine "Imaging data and services for aquatic science", which has received funding from the European Union Horizon Europe Programme –Grant Agreement number 101058625. Also, by the project AI4EOSC "Artificial Intelligence for the European Open Science Cloud", that has received funding from the European Union's HORIZON-INFRA-2021-EOSC-01 Programme under Grant 101058593. Finally, Grant PID2020-113126RB-I00 funded by MICIU/AEI/10.13039/501100011033.

Primary authors: RODRÍGUEZ, Vicente (Universitat Politècnica de València); CALATRAVA ARROYO, Amanda (Universitat Politècnica de València); AGUIRRE, Diego A. (Universitat Politècnica de València); LANGARITA, Sergio (Universitat Politècnica de València); ALARCÓN, Caterina (Universitat Politècnica de València); MOLTÓ, Germán (Universitat Politècnica de València)

Presenter: MOLTÓ, Germán (Universitat Politècnica de València)

Session Classification: IBERGRID

Track Classification: Development of innovative software and services