

OSCAR: Serverless Computing for Data-Processing Applications in the EGI Federated Cloud

Wednesday, 8 May 2019 09:25 (20 minutes)

Serverless computing, in the shape of an event-driven functions-as-a-service (FaaS) computing model, is being widely adopted for the execution of stateless functions that can be rapidly scaled in response to certain events. However, the main services offered by the public Clouds providers, such as AWS Lambda, Microsoft Azure Functions and Google Functions do not fit the requirements of scientific applications, which typically involve resource-intensive data processing and longer executions than those supported by the aforementioned services.

Still, scientific applications could benefit from the ability of being triggered in response to file uploads to a certain storage platform, so that the execution of multiple parallel invocations of the function/application would speed up the simultaneously processing of data while provisioning on-demand the required computing power to cope with the increased workloads.

Enter OSCAR, an open-source platform to support serverless computing for data-processing applications. OSCAR supports the FaaS computing model for file-processing applications. It can be automatically deployed on multi-clouds thanks to the EC3 (Elastic Cloud Compute Cluster) and IM (Infrastructure Manager) open-source developments.

OSCAR is provisioned on top of a Kubernetes cluster which is configured with a plugin created for the CLUES elasticity manager, in order to automatically provision additional nodes of the Kubernetes cluster to achieve two-level elasticity (elasticity for the number of containers and elasticity for the number of nodes). The following services are deployed inside the Kubernetes cluster: i) Minio, a high-performance distributed object storage server with an API compatible with Amazon S3; ii) OpenFaaS, a FaaS platform to create functions triggered via HTTP requests; iii) OSCAR UI, a web-based GUI aimed at end users to facilitate interaction with the platform.

The development of OSCAR has reached the status of prototype (TRL6) and it has also been integrated with the following EGI Services: i) EGI DataHub, in order to use OneData as the source of events. This way, scientists can upload files to their OneData space and this triggers invocations of the function in order to perform parallel processing on multiple files where the output data is automatically stored back in the same space; ii) EGI Applications on Demand, for users to self-deploy their elastic OSCAR clusters on EGI Federated Cloud through the EC3 portal; iii) EGI Cloud Compute, to provision Virtual Machines, as nodes of the Kubernetes cluster, from the EGI Federated Cloud.

The benefits of this platform have been assessed by integrating a use case related to Plants Classification using DEEP learning techniques that arised in the context of the DEEP Hybrid-DataCloud European project.

This activity is co-funded by the EGI Strategic and Innovation Fund.

Type of abstract

Presentation

References

OSCAR: <https://github.com/grycap/oscar>

Primary authors: Mr PÉREZ, Alfonso (Universitat Politècnica de València); Mrs NARANJO, Diana M. (Universitat Politècnica de València); Dr MOLTÓ, Germán (Universitat Politècnica de València); Dr CABALLER, Miguel (UPVLC); Mr RISCO, Sebastián (Universitat Politècnica de València)

Presenters: Mr PÉREZ, Alfonso (Universitat Politècnica de València); Dr MOLTÓ, Germán (Universitat Politècnica de València)

Session Classification: Cloud Platforms