

Cloud Orchestration at Application Level

Thursday, November 30, 2017 11:15 AM (15 minutes)

Overview. H2020 COLA (Cloud Orchestration at the Level of Application) project is developing the MiCADO (Microservices-based Cloud-application Level Orchestrator) platform. This platform supports the dynamic application level orchestration of cloud applications on multiple heterogeneous federated clouds. It enables execution of these applications based on their specific needs, such as cost, resources, security requirements, time constraints etc. in two phases: optimised deployment (1) and run-time orchestration (2).

In phase 1 application developers create a high-level description of their applications in TOSCA. This description, besides application topology, also includes various QoS (Quality of Service) parameters such as cost and performance requirements, and security policies. This description is passed on to the Coordination/Orchestration component of MiCADO. It collaborates with the Security facilitator that converts user-defined security policies into specific security solutions, and also with the Optimisation decision maker that translates cost and performance related parameters into actual deployment values. Following this, the Coordination/Orchestration component passes on the deployment instructions to MiCADO's Deployment executor that deploys the services required to run the application on the targeted cloud infrastructure. After deploying the application in the Cloud run-time orchestration starts (phase 2). MiCADO continuously collects various metrics from the running application and passes it on to the Coordination/orchestration component. This data is analysed from performance/cost aspects by the Optimisation decision maker and from security enforcement point of view by the Security facilitator. If adjustment is required, the Deployment executor is called to scale up or down the infrastructure. Users also have the possibility to adjust any requirement, either security or performance/cost related, during runtime. If the user provides a modified description with updated QoS parameters, then it is passed on to the Coordination/Orchestration component that analyses the received data and instructs the Deployment executor to modify the infrastructure.

MiCADO runs applications on a cluster to dynamically allocate and attach, or detach and release cloud resources for optimizing the resource usage. This cluster consists of two main components: Master node and Worker node. Master node is the head of the cluster performing the collection of information on microservices, the calculation of optimized resource usage, the decision making and the realization of decisions related to handling resources and to scheduling microservices. Worker nodes represent execution environments for microservices. Worker nodes are continuously allocated/released based on the dynamically changing requirements of the running microservices. Once a new worker node is allocated and attached to the cluster, the master node utilizes its resources by allocating microservices on it.

Conference themes and track topic. The presentation will address "thematic building blocks to the EOSC" topic (topic 5) outlining an orchestration approach that allows running applications on federated clouds. This approach enables application specific deployment and their scaling up/down considering user requirements.

Targeted audience. The presentation addresses application developers. They can describe application in TOSCA using a GUI, such as Alien4Cloud, deploy and run applications on federated clouds through the MiCADO platform.

Topic Area

The EOSC & EDI building blocks

Type of abstract

Presentation (15 minutes)

Primary author: TERSTYANSZKY, Gabor (University of Westminster)

Co-author: KISS, Tamas (University of Westminster, London, UK)

Presenter: TERSTYANSZKY, Gabor (University of Westminster)

Session Classification: EOSC building blocks presentations