

MiCADO - A highly customisable multi-cloud orchestration and auto-scaling framework

Wednesday, 8 May 2019 10:05 (20 minutes)

Many scientific and commercial applications require access to computation, data or networking resources based on dynamically changing requirements. Users and providers both require these applications or services to dynamically adjust to fluctuations in demand and serve end-users at a required quality of service (performance, reliability, security, etc.) and at optimized cost. This may require resources of these applications or services to automatically scale up or down.

The European funded, H2020 COLA (Cloud Orchestration at the Level of Application - <https://project-cola.eu/>) project set out to design and develop a generic framework that supports automated scalability of a large variety of applications. Learning from previous similar efforts and with the aim of reusing existing open source technologies wherever possible, COLA proposed a modular architecture called MiCADO (Microservices-based Cloud Application-level Dynamic Orchestrator - <https://www.micado-scale.eu/>) to provide optimized deployment and run-time orchestration for cloud applications.

MiCADO is open-source (<https://github.com/micado-scale>) and built from well-defined building blocks implemented as microservices. This modular design supports various implementations where components can be replaced relatively easily with alternative technologies. These building blocks, both on the MiCADO Master and also on the MiCADO Worker Nodes are implemented as microservices. The current implementation uses widely applied technologies, such as Kubernetes as the Container Orchestrator, Occopus as the Cloud Orchestrator, and Prometheus as the Monitoring System. When compared to similar auto-scaling frameworks, MiCADO is distinguished by its multi-cloud support, highly customisable scaling policies, policy-driven security settings, easy Ansible-based deployment, and its intuitive dashboard.

The user facing interface of MiCADO is a TOSCA-based (Topology and Orchestration Specification for Cloud Applications, an OASIS standard) Application Description Template which describes the desired container and virtual machine topology and its associated scalability and security policies. This interface has the potential to be embedded to existing GUIs, custom web interfaces or science gateways.

MiCADO has been tested with several large-scale industry and research applications, and on various public (AWS, CloudSigma, MS Azure, CloudBroker) and private (OpenStack, OpenNebula) cloud resources. The two main targeted application types are cloud-based services where scalability is achieved by scaling up or down the number of containers and virtual machines based on load, performance and cost, and the execution of a large number of jobs where a particular experiment consisting of hundreds or thousands of jobs needs to be executed by a set deadline.

The proposed presentation will give an overview of MiCADO, will explain its architecture and characteristics, and will demonstrate via application case-studies how EGI user communities could utilise this technology.

Primary author: Prof. KISS, Tamas (University of Westminster)

Co-authors: Mr OCKLENBURG, Andreas (CloudSME); Prof. TERSTYANSZKY, Gabor (University of Westminster); Dr PIERANTONI, Gabriele (University of Westminster); Mr GESMIER, Gregoire (University of Westminster); Mr DESLAURIERS, James (University of Westminster); Dr KOVACS, Jozsef (MTA SZTAKI); Prof. KACSUK, Peter (MTA SZTAKI)

Presenter: Prof. KISS, Tamas (University of Westminster)

Session Classification: Cloud Platforms