

Infrastructure Manager: A decade of innovation in virtualised infrastructure provisioning

Tuesday, 26 September 2023 13:00 (20 minutes)

Scientific computing benefits from the automated provisioning of virtualized infrastructures from multiple Infrastructure as a Service (IaaS) clouds. An abstraction on how computational resources are provisioned, configured and delivered to the end user is required to widespread the adoption of Cloud computing. During the last decade, the development of the open-source Infrastructure Manager (IM) tool has provided the ability to interact with multiple on-premises Cloud Management Platforms (CMPs), such as OpenStack and OpenNebula, many public Cloud providers, such as Amazon Web Services (AWS), Microsoft Azure or Google Cloud Platform, including European public Cloud providers such as Cloud & Heat, Open Telekom Cloud or Orange Cloud and also Federated providers as EGI Cloud Compute or FogBow.

The IM leverages open-source technologies for Infrastructure as Code (IaC) such as Ansible and standards such as TOSCA (Topology and Open Specification for Cloud Applications), to describe complex application architectures to be deployed on the cloud. The IM supports multiple integration paths for different user profiles, including a fully-featured REST API, alongside an easy-to-use command-line interface. In addition, the web-based IM Dashboard facilitates the usage by less savvy users by offering pre-packaged tested popular application architectures that can seamlessly be deployed on multiple Cloud back-ends. This is the case of Kubernetes clusters, JupyterLab instances, Hadoop clusters, etc. By providing a wizard-like TOSCA-based composition approach, users can further customize their deployment specifying the resources to allocate, the software configuration or the number of nodes to deploy in the cluster. It also provides the ability to scale the deployed infrastructures both horizontally (adding or removing nodes) or vertically (resizing a particular VM).

This contribution looks back into a decade of innovation in the field of automated provisioning of virtualised infrastructures using the Infrastructure Manager, highlighting its usage in multiple European projects (e.g. INDIGO-DataCloud, EOSC-HUB, EGI-ACE, AI-SPRINT, InterTwin, DT-GEO, etc.) and scientific communities (e.g. PanGeo, ENES, etc.). As a result, the Infrastructure Manager is being offered in production as one of the EGI services for research.

This contribution will also address the evolution of the Infrastructure Manager to become an orchestration component for the edge-to-cloud continuum, by allowing the definition of event-driven functions to be deployed on on-premises serverless computing platforms such as OSCAR and public FaaS offerings such as AWS Lambda.

References

- [1] Infrastructure Manager. <https://im.egi.eu>
- [2] Infrastructure Manager. <https://www.egi.eu/service/infrastructure-manager/>
- [3] OSCAR. <https://oscar.grycap.net>

Primary authors: CALATRAVA ARROYO, Amanda (Universitat Politècnica de València); Dr MOLTÓ, Germán (Universitat Politècnica de València); CABALLER, Miguel (Universitat Politècnica de València)

Presenter: CABALLER, Miguel (Universitat Politècnica de València)

Session Classification: IBERGRID Contributions with Demonstrations

Track Classification: Development of innovative software and services