

Virtualized Web Portals in EGI Federated Cloud

Web portals became a popular interface to complicated workflows of scientific computation in many areas due to their ability to shield the user from the complexity of the workflow implementation and control of the used computing resources.

On the other hand, implementation of such web portal itself involves rather complicated technology, which becomes difficult to operate and maintain as the number of its users grows. Moreover, user requirements may become contradictory (e.g. stick with an old version of simulation software to retain full backward compatibility vs. use a new version to leverage new features).

Cloud infrastructures available nowadays, together with cloud orchestration engines allow addressing those issues in a scalable way. Rather than maintaining “one size fits all” instance of an application web portal, we provide templates for deployment multiple instances of a specific portal, dedicated to smaller user groups. Such instances may vary in size, and they can be easily customized to the needs of their users, while still reusing 99% of the common code and configuration.

We describe in detail the solution adopted in the EU H2020 West-life project. The current implementation uses Cloudify orchestrator, extended with OCCI plugin, to spawn and to keep alive the set of virtual machines forming the portal in the environment of EGI Federated Cloud. Deliberately, minimalistic VM images containing the bare operating system only are used; all required software is installed and configured with Puppet system management upon VM instantiation. In this way, tedious management of huge VM images, which would require frequent updates and which would be complicated to distribute otherwise, is avoided.

We demonstrate the approach with a prototype of web portal on top of the Gromacs molecular dynamics package. The original version of the portal was developed in the WeNMR EU project, it is based on Gromacs version 4.5, and it uses gLite workload management system to submit payload jobs to the EGI grid resources. Besides wrapping the portal to the virtualized cloud environment we adapted it to use the more recent version of Gromacs (5.1), to leverage GPU acceleration when the hardware is available. Submission to the grid was replaced with use of local batch system (Torque) which is sufficient to manage resources at the size expected for any instance of such virtualized portal, while being less complicated, self-contained, and more reliable.

We also outline plans for future work – adopting the more advanced cloud orchestrator developed in EU Indigo-Datacloud project; in particular, we will focus on flexible sizing of the portal instances according to their current load.

Primary author: Dr KŘENEK, Aleš (Masaryk University)

Presenter: Dr KŘENEK, Aleš (Masaryk University)

Track Classification: Virtual Research Environment (including Middleware, tools, services, workflow, ... etc.)