

Batch virtualization project at CERN

Monday 26 October 2009 15:30 (30 minutes)

Between March and August 2009 a project has been set up at CERN with the aim to evaluate possibilities to use virtualization at a large scale, with the focus on batch computing. Two key issues have been identified for this specific application: the placing of virtual machines on an appropriate hyper-visor, and the selection of an appropriate image which should be driven by the actual demand. Both commercial and free software solutions exist which are able to solve the placing issue.

The virtual machine orchestrator, VMO, a commercial solution by Platform computing, and the free software solution OpenNebula have been evaluated during the project. For VMO, the vendor provided a first implementation of an algorithm for selecting the image to be deployed, which is driven by user requirements of pending batch jobs. For OpenNebula, an external mechanism needs to be developed to perform this task. In the presentation, the basic concepts of the project and lessons learned will be presented. Further visions and possible implications for services offered at CERN will be described.

Primary author: Dr SCHWICKERATH, Ulrich (CERN)

Co-authors: Dr GOASGUEN, Sebastien (Clemson University); Dr CASS, Tony (CERN)

Presenter: Dr SCHWICKERATH, Ulrich (CERN)

Session Classification: Virtualization I

Track Classification: Virtualization