

Use Case: Running Monte Carlo LHCb simulations using DIRAC with EGI Federated Cloud

Thursday, 20 September 2012 11:00 (20 minutes)

Description of the work

DIRAC is integrated with private, public and hybrid clouds. In combination with CernVM, it provides a very flexible platform for general purpose usage and for this reason it has been chosen for this use case. Minor developments have been required in order to allow multiple cloud providers to be managed from a single DIRAC instance.

The resulted extension to the DIRAC Virtual Engine needs to be tested before being used in production. After performing some small-scale functionality tests, the goal is to execute a large-scale test aiming to consume 2000 core hours at each site on the EGI Federated Cloud infrastructure. Small Virtual Machines (VM) of 1 Core and 2 GB of memory are required. The deployed images are the CernVM with about 14GB of filesystem size on local disk. The scaling test will be run in multiple clouds at the same time.

After the integration and scaling testing, the cloud resources will be connected to the LHCb DIRAC production instance and configure the VMs to match and execute MC simulation jobs. Site Managers will establish a maximum number of virtual machine images running at the same time. The expected contribution from each site, integrated over the duration of this last phase, should be of the order of several thousands of core days. The LHCb Production Managers will take care of the production from LHCb-DIRAC portal and are not supposed to see any difference beyond the additional computing resources available.

The possibility to connect commercial public CPU providers with a contribution at the same level will also be explored.

Link for further information

<http://diracgrid.org>

Wider impact of this work

The use case will prove the viability of a Monte Carlo production using Grid and Cloud resources through DIRAC. This is one of the DIRAC team efforts to provide wider integration of resources, users and tools making DIRAC “the interware” solution. At the same time this use case is contributing to the EGI Federation Cloud aim of a production level.

The expertise of this use case will be useful for other LHCb activities. Furthermore, the experience can be directly transferred to other communities already using DIRAC interware solution (ILC, Belle II, CTA, BES III), and to the NGI infrastructures already providing DIRAC interware as Service for their national user communities (France-Grilles and IberGrid).

Printable Summary

The DIRAC project provides a framework for building ready to use distributed computing systems. It has been proven to be a useful tool for large international scientific collaborations integrating in a single system, their computing activities and distributed computing resources: Grids, Clouds and HTC clusters. In the case of Cloud resources, DIRAC is currently integrated with Amazon EC2, OCCI/OpenNebula and CloudStack. Some Monte Carlo (MC) simulation campaign were realized at the large scale project Belle II, providing over 10.000 thousand CPU days from Amazon. Until now, all cases have made used of a single cloud at a time.

The work aims to integrate the resources provided by the multiple private clouds of the EGI Federated Cloud and additional WLCG resources, providing high-level scientific services on top of them by using the DIRAC framework. New design and development are discussed. Initial integration and scaling tests are necessary to anticipate possible problems.

Primary author: Dr MENDEZ, Victor (Port d'Informació Científica, PIC)

Co-authors: Mr ADRIÀ, Casajús (Universidad de Barcelona); Dr MERINO, Gonzalo (Port d'Informació Científica, PIC); Dr DIAZ, Ricardo (Universidad de Barcelona)

Presenter: Dr MENDEZ, Victor (Port d'Informació Científica, PIC)

Session Classification: Virtualised Resources

Track Classification: Virtualised Resources: challenges and opportunities (Michel Drescher: track leader)