



Contribution ID: 94

Type: **Parallel**

CernVM Co-Pilot: an Extensible Framework for Building Scalable Cloud Computing Infrastructures

Tuesday 22 May 2012 13:30 (25 minutes)

CernVM Co-Pilot is a framework for instantiating an ad-hoc computing infrastructure on top of distributed computing resources. Such resources include commercial computing clouds (e.g. Amazon EC2), scientific computing clouds (e.g. CERN lxcloud), as well as the machines of users participating in volunteer computing projects (e.g. BOINC). The framework consists of components that communicate using the Extensible Messaging and Presence protocol (XMPP), allowing for new components to be developed in virtually any programming language and interfaced to existing Grid and batch computing infrastructures exploited by the High Energy Physics community. Co-Pilot has been used to execute jobs for both the ALICE and ATLAS experiments at CERN.

CernVM Co-Pilot is also one of the enabling technologies behind the LHC@home 2.0 volunteer computing project, which is the first such project that exploits virtual machine technology. The use of virtual machines eliminates the necessity of modifying existing applications and adapting them to the volunteer computing environment. After start of the public testing in August 2011 LHC@home 2.0 quickly gained popularity, and as of October 2011 it had about 9000 registered volunteers. Resources provided by volunteers are used for running Monte-Carlo generator applications that simulate interactions between the colliding proton beams at the LHC.

In this contribution we present the latest developments and the current status of the system, discuss how the framework can be extended to suit the needs of a particular scientific community, describe the operational experience using the LHC@home 2.0 volunteer computing infrastructure, as well as introduce future development plans.

Summary

CernVM Co-Pilot is a framework for instantiating an ad-hoc computing infrastructure on top of academic and commercial computing clouds, or on the machines of users participating in volunteer computing projects. The framework consists of components that communicate using the Extensible Messaging and Presence protocol (XMPP), allowing for new components to be developed in virtually any programming language and interfaced to existing Grid and batch computing infrastructures.

In this contribution we present the latest developments and the current status of the system, discuss how the framework can be extended to suit the needs of a particular scientific community, describe the operational experience using the LHC@home 2.0 volunteer computing infrastructure, as well as introduce future development plans.

Author: HARUTYUNYAN, Artem (CERN)

Co-authors: LARSEN, Dag (University of Bergen (NO)); CHARALAMPIDIS, Ioannis (Aristotle Univ. of Thessaloniki (GR)); BLOMER, Jakob (Ludwig-Maximilians-Univ. Muenchen (DE)); BUNCIC, Predrag (CERN)

Presenter: HARUTYUNYAN, Artem (CERN)

Session Classification: Distributed Processing and Analysis on Grids and Clouds

Track Classification: Distributed Processing and Analysis on Grids and Clouds (track 3)