



Contribution ID: 45

Type: **Poster presentation**

Cloud flexibility using Dirac Interware

Monday, 14 October 2013 15:00 (45 minutes)

Communities of different locations are running their computing jobs on dedicated infrastructures without the need to worry about software, hardware or even the site where their programs are going to be executed. Nevertheless, this usually implies that they are restricted to use certain types or versions of an Operating System because either their software needs a definite version of a system library or a specific platform is required by the collaboration to which they belong. On this scenario, if a data center wants to service software incompatible communities, it has to split its physical resources among those communities. This splitting will inevitably lead to an underuse of resources because the data centers are bound to have periods where one or more of its subclusters are idle.

It is in this situation where Cloud Computing provides the flexibility and reduction in computational cost that data centers are searching for. This paper describes a set of realistic tests that we ran on one of such implementations. The test comprise software from three different HEP communities (Auger, LHCb and QCD phenomenologists) and the Parsec Benchmark Suite running on one or more of three Linux flavors (SL5, Ubuntu 10.04 and Fedora 13). The implemented infrastructure has, at the cloud level, CloudStack that manages the virtual machines (VM) and the hosts on which they run, and, at the user level, the DIRAC framework along with a VM extension that will submit, monitorize and keep track of the user jobs and also requests CloudStack to start or stop the necessary VM's. In this infrastructure, the community software is distributed via the CernVM-FS, which has been proven to be a reliable and scalable software distribution system. With the resulting infrastructure, users are allowed to send their jobs transparently to the Data Center.

The main purpose of this system is the creation of flexible cluster, multiplatform with an scalable method for software distribution for several VOs. Users from different communities do not need to care about the installation of the standard software that is available at the nodes, nor the operating system of the host machine, which is transparent to the user.

Summary

Primary authors: SECO MIGUELEZ, Marcos (Universidade de Santiago de Compostela (ES)); FERNANDEZ ALBOR, Victor Manuel (Universidade de Santiago de Compostela (ES))

Co-authors: SABORIDO SILVA, Juan Jose (Universidade de Santiago de Compostela (ES)); GRACIANI DIAZ, Ricardo (University of Barcelona (ES)); Prof. FERNANDEZ, Tomas (Universidade de Santiago de Compostela (ES)); MENDEZ MUÑOZ, Victor (Port d'Informació Científica (PIC), Universitat Autònoma de Barcelona, ES-08193 Bellaterra (Barcelona) Spain)

Presenter: FERNANDEZ ALBOR, Victor Manuel (Universidade de Santiago de Compostela (ES))

Session Classification: Poster presentations

Track Classification: Distributed Processing and Data Handling A: Infrastructure, Sites, and Virtualization