

Exploring virtualization tools with a new virtualization provisioning method to test dynamic grid environments for ALICE grid jobs over ARC grid middleware

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

The Nordic Tier-1 for the LHC is distributed over several, sometimes smaller, computing centers. In order to minimize administration effort, we are interested in running different grid jobs over one common grid middleware. ARC is selected as the internal middleware in the Nordic Tier-1. The AliEn grid middleware, used by ALICE has a different design philosophy than ARC. In order to use most of the AliEn infrastructure and available software deployment methods for running ALICE grid jobs on ARC, we are investigating different possible virtualization technologies.

Software deployment is a critical a part in grid computing. Ideally a contributing computing center doesn't need to change their infrastructure or install additional software, except at the grid entry point. One solution to achieve this is a cloud infrastructure that provides (dynamically) virtual machines with the required software environments. It seems that no easy solution exists to run grid jobs on top of those cloud providers. They are missing tools that submit a job into a virtual machine, that may be created dynamically and collect the results afterwards. Therefore a specialized, virtualization provisioning method was developed. This method could be developed further into a system that is used by smaller sites, which don't want to run a more general and more complex private cloud solution.

The new provisioning system is used to test different combinations of virtualization backends and virtual machines.

The CernVM project is developing a virtual machine that can provide a common analysis environment for all LHC experiments. One of our interests is to investigate the use of CernVM as a base setup for a dynamical grid environment capable of running grid jobs. For this, performance comparisons between different virtualization technologies have been conducted.

CernVM needs an existing virtualization infrastructure, which is not always existing or wanted at some computing sites. To increase the possible application of dynamical grid environments to those sites, we describe several possibilities that are less invasive and have less specific Linux distribution requirements, at the cost of lower performance.

Different tools like user-mode Linux (UML), micro Linux distributions, a new software packaging project by Stanford university (CDE) and CernVM are under investigation for their invasiveness, distribution requirements and performance. Comparisons between the different methods with solutions that are closer to the hardware will be presented.

Summary

Primary authors: KILENG, Bjarte (Bergen University College (NO)); WAGNER, Boris (University of Bergen (NO) for the ALICE Collaboration)

Presenter: WAGNER, Boris (University of Bergen (NO) for the ALICE Collaboration)

Session Classification: Poster presentations

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