Contribution ID: **504** Type: **Parallel**

Connecting multiple clouds and mixing real and virtual resources via the open source WNoDeS framework

Tuesday, May 22, 2012 3:10 PM (25 minutes)

In this paper we present the latest developments introduced in the WNoDeS framework (http://web.infn.it/wnodes); we will in particular describe inter-cloud connectivity, support for multiple batch systems, and coexistence of virtual and real environments on a single hardware.

Specific effort has been dedicated to the work needed to deploy a "multi-sites" WNoDeS installation. The goal is to give end users the possibility to submit requests for resources using cloud interfaces on several sites in a transparent way. To this extent, we will show how we have exploited already existing and deployed middleware within the framework of the IGI (Italian Grid Initiative) and EGI (European Grid Infrastructure) services. In this context, we will also describe the developments that have taken place in order to have the possibility to dynamically exploit public cloud services like Amazon EC2. The latter gives WNoDeS the capability to serve, for example, part of the user requests through external computing resources when needed, so that peak requests can be honored.

We will then describe the work done to add support for open source batch systems like Torque/Maui to WNoDeS. This makes WNoDeS a fully open source product and gives the possibility to smaller sites as well (where often there is no possibility to run commercial batch systems) to install it and exploit its features. We will also describe recent WNoDeS I/O optimizations, showing results of performance tests executed using Torque as batch system and Lustre as a distributed file system.

Finally, starting from the consideration that not all tasks are equally suited to run on virtual environments, we will describe a novel feature added to WNoDeS, allowing the possibility to use the same hardware to run both virtual machines and real jobs (i.e., jobs running on the bare metal and not in a virtualized environment). This increases flexibility and may optimize the usage of available resources. In particular, we will describe tests performed in order to show how this feature can help in fulfilling requests for "whole-node jobs" (which are becoming increasingly popular in the HEP community), for efficient analysis support, and for GPU-based resources (which are typically not easily amenable to be virtualized).

Primary authors: Mr ITALIANO, Alessandro (INFN-CNAF); CHIERICI, Andrea (INFN-CNAF); ANDREOTTI, Daniele (Universita e INFN (IT)); SALOMONI, Davide (Istituto Nazionale Fisica Nucleare (IT)); RONCHIERI, Elisabetta (Universita e INFN (IT)); Dr DONVITO, Giacinto (INFN-Bari); DALLA TORRE, Gianni; SPINOSO, Vincenzo (Universita e INFN (IT))

Presenters: Mr ITALIANO, Alessandro (INFN-CNAF); Dr DONVITO, Giacinto (INFN-Bari)

Session Classification: Distributed Processing and Analysis on Grids and Clouds

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