



Contribution ID: 118

Type: **Poster presentation**

A Scalable Infrastructure for CMS Data Analysis Based on Openstack Cloud and Gluster File System

Monday, October 14, 2013 3:00 PM (45 minutes)

The challenge of providing a resilient and scalable computational and data management solution for massive scale research environments, such as the CERN HEP analyses, requires continuous exploration of new technologies and techniques. In this article we present a hybrid solution of an open source cloud with a network file system for CMS data analysis. Our aim has been to design a scalable and resilient infrastructure for CERN HEP data analysis. The infrastructure is based on Openstack components for structuring a private cloud together with the Gluster filesystem. The Openstack cloud platform is one of the fastest growing solutions in the cloud world. The Openstack components provide solutions for computational resource (NOVA), instance (GLANCE), network (QUANTUM) and security (Keystone) management all underneath an API layer that supports global applications, web clients and large ecosystems. NOVA and GLANCE manage Virtual Machines (VMs) and image repository respectively. QUANTUM provides network as a service (NAAS) and advanced network management capabilities. The virtual network layer provided by QUANTUM supports seamless migration of the VMs while preserving the network configuration. One important component that is currently not part of the Openstack suite is a network file system. To overcome this limitation we have used GlusterFS, a network-based file system for high availability. GlusterFS uses FUSE and may be scale to petabytes. In our experiments, 1TB is used for instance management and 2TB for the data related to CMS jobs within GlusterFS. We integrate the above state-of-the-art cloud technologies with the traditional Grid middleware infrastructure. This approach implies no changes for the end-user, while the production infrastructure is enriched by the high-end resilient and scalable components. To achieve this, we have run Advance Resource Connector (ARC) as a meta-scheduler. Both Computing Elements (CE) and Worker Nodes (WN) are running on VM instances inside the Openstack cloud. Currently we consider our approach as semi-static, as the instance management is manual yet provides scalability and performance. In near future we are aiming for a comprehensive elastic solution by including the EMI authorization service (Argus) and the Execution Environment Service (Argus-EES). In order to evaluate the strength of the infrastructure, four test cases have been selected for experimentation and analysis. (i) The first test case is based on instance performance, the boot time of customized images using different hypervisors and the performance of multiple instances with different configurations. (ii) The second test case focuses on I/O-related performance analysis based on GlusterFS. This test also presents the performance of instances running on GlusterFS compared with the local file system. (iii) The third test case examines system stability with live migration of VM instances based on QUANTUM. (iv) In the fourth test we will present long-term system performance both at the level of VMs running CMS jobs and physical hosts running VMs. Our test results show that the adopted approach provides a scalable and resilient solution for managing resources without compromising on performance and high availability.

Primary authors: WHITE, John White (Helsinki Institute of Physics (FI)); OSMANI, Lirim (University of Helsinki); KRAEMER, Oscar (Helsinki Institute of Physics (HIP)); EEROLA, Paula (Helsinki Institute of Physics (HIP)); Dr TOOR, Salman (Helsinki Institute of Physics (FI)); TARKOMA, Sasu (University of Helsinki); LINDÉN, Tomas (HELSINKI INSTITUTE OF PHYSICS)

Presenter: Dr TOOR, Salman (Helsinki Institute of Physics (FI))

Session Classification: Poster presentations

Track Classification: Facilities, Production Infrastructures, Networking and Collaborative Tools