20th International Conference on Computing in High Energy and Nuclear Physics (CHEP2013)

Contribution ID: **481** Type: **Oral presentation to parallel session**

Implementation of grid Tier 2 and Tier 3 facilities on a Distributed OpenStack Cloud

Monday, 14 October 2013 16:51 (22 minutes)

The Australian Government is making a $AUD 100 million investment in Compute and Storage for the academic community. The Compute facilities are provided in the form of 24,000 CPU cores located at 8 nodes around Australia in a distributed virtualized Infrastructure as a Service facility based on OpenStack. The storage will eventually consist of over 100 petabytes located at 6 nodes. All will be linked via a 100 Gbs network.

This presentation will describe the development of a fully connected WLCG Tier-2 grid site as well as a general purpose Tier-3 computing cluster based on this architecture.

The facility employs an extension to Torque to enable dynamic allocations of virtual machine instances. Storage is provided by a federation of DPM installations at each storage node. A base Scientific Linux VM image is deployed in the OpenStack cloud and automatically configured as required using Puppet. Custom scripts are used to launch multiple VMs, integrate them into the dynamic Torque cluster and to mount remote file systems.

We will report on our experience in implementing this nation-wide ATLAS and Belle II Tier 2 and Tier 3 computing infrastructure using the
national Research Cloud and storage facilities. In particular we will describe how we have addressed the challenges of using OpenStack VMs in a Torque cluster, automated configuration of VM instances, federated authentication across multiple institutions and supported access to
remote file systems.

Summary

Primary author: SEVIOR, Martin (University of Melbourne)

Co-authors: Dr LIMOSANI, Antonio (University of Melbourne (AU)); Ms HUANG, Joanna (University of Melbourne); Mr WILSON, Ross (University of Adelaide); ZHANG, Shunde (eRSA)

Presenter: Dr LIMOSANI, Antonio (University of Melbourne (AU))

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

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