



Contribution ID: 163

Type: Oral

Dynamic provisioning of a HEP computing infrastructure on a shared hybrid HPC system

Monday 18 January 2016 15:45 (25 minutes)

The Institut für Experimentelle Kernphysik (EKP) at KIT is a member of the CMS and Belle II experiments, located at the LHC and the Super-KEKB accelerators, respectively. These detectors share the requirement, that enormous amounts of measurement data must be processed and analyzed and a comparable amount of simulated events is required to compare experimental results with theory predictions.

Nowadays, funding agencies encourage research groups to participate in shared HPC cluster models, where scientist from different domains use the same hardware to increase synergies. This shared usage proves to be challenging for high-energy physics (HEP) groups, due to their specialized software setup which includes a custom OS (often Scientific Linux), libraries and applications.

To overcome this hurdle, the EKP and data center team of the University of Freiburg have developed a system to enable the HEP use case on a shared HPC cluster. To achieve this, an OpenStack-based virtualization layer is installed on-top of a bare-metal cluster. While other user groups can run their batch jobs via the Moab workload manager directly on bare-metal, HEP users can request virtual machines with a specialized machine image which contains a dedicated operating system and software stack. Contrary to similar installations, in this hybrid setup, no static partitioning of the cluster into a physical and virtualized segment is required.

A seamless integration with the jobs sent by other users groups honors the fairshare policies of the cluster. The developed thin integration layer between OpenStack and Moab can be adapted to other batch servers and virtualization systems, making the concept also applicable for other cluster operators.

This contribution will report on the concept and implementation of an OpenStack-virtualized cluster used for HEP workflows. While the full cluster will be installed in spring 2016, a test-bed setup with 800 cores has been used to study the overall system performance and dedicated HEP jobs were run in a virtualized environment over many weeks. Furthermore, the dynamic integration of the virtualized worker nodes, depending on the workload at the institute's computing system, will be described.

Authors: FLEIG, Georg Stefan (KIT - Karlsruhe Institute of Technology (DE)); MEIER, Konrad (Albert-Ludwigs-Universität Freiburg); HAUTH, Thomas (KIT)

Co-authors: WIEBELT, Bernd (Albert-Ludwigs-Universität Freiburg); VON SUCHODOLETZ, Dirk (Albert-Ludwigs-Universität Freiburg); QUAST, Gunter (KIT - Karlsruhe Institute of Technology (DE)); JANCZYK, Michael (Albert-Ludwigs-Universität Freiburg)

Presenter: MEIER, Konrad (Albert-Ludwigs-Universität Freiburg)

Session Classification: Track 1

Track Classification: Computing Technology for Physics Research