



Contribution ID: 39

Type: **not specified**

Fermilab HEP Cloud: an elastic computing facility for High Energy Physics.

Tuesday, 19 April 2016 11:50 (25 minutes)

The need for computing in the HEP community follows cycles of peaks and valleys mainly driven by holiday schedules, conference dates and other factors. Because of this, the classical method of provisioning these resources at providing facilities has drawbacks such as potential overprovisioning. As the appetite for computing increases, however, so does the need to maximize cost efficiency by developing a model for dynamically provisioning resources only when needed.

To address this issue, the HEP Cloud project was launched by the Fermilab Scientific Computing Division at Fermilab in June 2015. Its goal is to develop a facility that provides a common interface to a variety of resources, including local clusters, grids, high performance computers, and community and commercial Clouds.

Initially targeted communities include CMS and NOvA, as well as other Fermilab stakeholders. In its first phase, the project has demonstrated the use of the “elastic” provisioning model offered by commercial clouds, such as Amazon Web Services. In this model, resources are rented and provisioned automatically over the Internet upon request. Cost was contained by the use of the Amazon Spot Instance Market, a rental model that allows Amazon to sell their overprovisioned capacity at a fraction of the regular price. Data access was made to scale in terms of volume and cost through a variety of techniques, including autoscaling data caching services.

In January 2016, the project demonstrated the ability to increase the total amount of global CMS resources by 58,000 cores from 150,000 cores - a 25 percent increase. This burst of resources was used in preparation for the Recontres de Moriond conference to generate and reconstruct Monte Carlo events. At the same time, the NOvA experiment has also run data-intensive computations through HEP Cloud, readily provisioning 7,500 cores on Amazon to process Monte Carlo and reconstructed detector data for Neutrino conferences. NOvA is using the same familiar services they use for local computations such as data handling and job submission.

This talk will discuss the architecture used, lessons learned along the way, and some of the next steps in the evolution of the Fermilab HEP Cloud Facility.

Length of presentation (minutes, max. 20)

20

Primary author: TIRADANI, Anthony (Fermilab)

Presenter: TIRADANI, Anthony (Fermilab)

Session Classification: Grids, clouds, virtualisation

Track Classification: Grid, Cloud & Virtualisation