



Contribution ID: 8

Type: **not specified**

FermiCloud On-demand Services: Data-Intensive Computing on Public and Private Clouds

Friday, May 23, 2014 9:25 AM (25 minutes)

The FermiCloud project exists to provide on-demand computing and data movement services to the various experiments at Fermilab. We face a dynamically changing demand for compute resources and data movement, which we meet by enabling users to run on our own site, remote grid sites, and cloud sites. We also instantiate on-demand data movement and web caching services to support this remote analysis. In this presentation we will summarize some of our recent research results and outline the challenges of our current research projects. These include coordinated launches of compute nodes and data movement servers, interoperability with new commercial clouds, idle machine detection, and exploration of distributed storage models.

Summary

We have successfully demonstrated several methods of launching compute nodes to the public cloud including the vcluster software, the GlideinWMS, and OpenNebula cloud bursting features. Using the GlideinWMS we have submitted several workflows from the NOvA experiment to FermiCloud and the public Amazon EC2 cloud. Current research and discussion centers on tiered policy-based provisioning. A large input to the provisioning model is the cost in time and bandwidth that it takes to transfer virtual machine images and launch them, which we have modeled in detail. We have also published results on high-performance calculations using virtualized Infiniband adapters.

Current research is focused on coordinated provisioning of diverse cloud-based scientific workflow. We are also developing an automated conversion service to change between different virtual machine image formats. We are developing a set of tools to port applications to run both on external grid sites and on public clouds. We will summarize the best practices that we have found in running physics applications on public cloud as well as investigations of running using spot pricing. Finally we continue to evaluate other commercial and community clouds.

Author: TIMM, Steven (Fermilab)

Co-author: Dr GARZOGGIO, Gabriele (FERMI NATIONAL ACCELERATOR LABORATORY)

Presenter: TIMM, Steven (Fermilab)

Session Classification: Grids, clouds, virtualisation

Track Classification: Grid, Cloud & Virtualisation