Contribution ID: 59

Successful Running of Fermilab Experimental Simulation on Federated Grids and Clouds

Tuesday, 20 May 2014 14:40 (20 minutes)

Fermilab neutrino experimental analysis and simulation was required to run mostly at Fermilab due to access to a NAS server for their code and data. Fermilab staff worked with the experimenters to move their application to the CVMFS distribution and then run their application on several distributed grid sites, on FermiCloud, and on Amazon AWS. We found both internal and commercial cloud service to give performance and reliability similar to an opportunistic Open Science Grid site. We will present comparative execution times and efficiencies, work that was done, and suggested best practices for porting and running applications on external grids and clouds.

Wider impact and conclusions

Several other neutrino and muon experiments at Fermilab are now following the same road map to make their codes available to the grid and cloud. We expect to be able to get significant extra computing cycles on the grid, and are in the planning phases of a production batch facility which takes virtual machines submitted as batch jobs.

Description of work

The NOvA Monte Carlo application was commissioned to run on distributed resources by focusing on the software distribution capabilities, the access to remote resources through standard Grid a Cloud interfaces, and the output transfer mechanism. In particular, the code was deployed on the CernVM File System (CVMFS) repositories at Fermilab and in the OASIS service of the Open Science Grid (OSG) for on-demand distribution to all remote resources. A resource pool was commissioned with seven OSG sites, three dedicated and four opportunistic, as well as FermiCloud. The output data transfer mechanism was set up to cache the results at the remote site and then spool them back in bulk. The NOvA production team operated the system to run their Monte Carlo campaign and have continued to use it since.

Primary author: TIMM, Steven (Fermi National Accelerator Laboratory)

Co-authors: NORMAN, Andrew (Fermilab); DYKSTRA, Dave (Fermi National Accelerator Laboratory); BOX, Dennis (Fermilab); GARZOGLIO, Gabriele (Fermilab); KIM, Hyun Woo (Fermilab); BOYD, Joseph (Fermilab); MEN-GEL, Marc (Fermilab); KIRBY, Michael (Fermilab); MHASHILKAR, Parag (Fermilab); LEVSHINA, Tanya (Fermilab/OSG)

Presenter: DYKSTRA, Dave (Fermi National Accelerator Laboratory)

Session Classification: EGI Federated cloud: state of the art and future steps

Track Classification: Advanced cloud services (Track Leaders: D. Wallom, M. Drescher, K. Nordlund)