

Virtual Machine Provisioning, Code Management and Data Movement Design for the Fermilab HEPCloud Facility

Thursday 13 October 2016 16:30 (15 minutes)

The Fermilab HEPCloud Facility Project has as its goal to extend the current Fermilab facility interface to provide transparent access to disparate resources including commercial and community clouds, grid federations, and HPC centers. This facility enables experiments to perform the full spectrum of computing tasks, including data-intensive simulation and reconstruction. We have evaluated the use of the commercial cloud to provide elasticity to respond to peaks of demand without overprovisioning local resources. Full scale data-intensive workflows have been successfully completed on Amazon Web Services for two High Energy Physics Experiments, CMS and NOvA, at the scale of 58000 simultaneous cores. This paper describes the significant improvements that were made to the virtual machine provisioning system, code caching system, and data movement system to accomplish this work. The virtual image provisioning and contextualization service was extended to multiple AWS regions, and to support experiment-specific data configurations. A prototype Decision Engine was written to determine the optimal availability zone and instance type to run on, minimizing cost and job interruptions. We have deployed a scalable on-demand caching service to deliver code and database information to jobs running on the commercial cloud. It uses the frontier-squid server and CERN VM File System (CVMFS) clients on EC2 instances and utilizes various services provided by AWS to build the infrastructure (stack). We discuss the architecture and load testing benchmarks on the squid servers. We also describe various approaches that were evaluated to transport experimental data to and from the cloud, and the optimal solutions that were used for the bulk of the data transport. Finally we summarize lessons learned from this scale test, and our future plans to expand and improve the Fermilab HEP Cloud Facility.

Tertiary Keyword (Optional)

Distributed data handling

Secondary Keyword (Optional)

Computing facilities

Primary Keyword (Mandatory)

Cloud technologies

Primary author: TIMM, Steven (Fermilab)

Co-authors: TIRADANI, Anthony (Fermilab); HOLZMAN, Burt (Fermi National Accelerator Lab. (US)); Mr GRASSANO, Davide (Fermilab); GARZOGLIO, Gabriele; Mr WU, Hao (Illinois Institute of Technology); Dr KIM, Hyun Woo (Fermilab); Mr COOPER, R. Glenn (Fermilab); Mr KRISHNAMURTHY, Rahul (Illinois Institute of Technology); KENNEDY, Robert (Fermilab); Prof. REN, Shangping (Illinois Institute of Technology); Mr VINAYAGAM, Shivakumar (Illinois Institute of Technology); FUESS, Stuart (Fermilab)

Presenters: HOLZMAN, Burt (Fermi National Accelerator Lab. (US)); GARZOGLIO, Gabriele; TIMM, Steven (Fermilab); FUESS, Stuart (Fermilab)

Session Classification: Posters B / Break

Track Classification: Track 3: Distributed Computing