

Advances in Grid Computing for the Fabric for Frontier Experiments Project at Fermilab

Tuesday, October 11, 2016 11:45 AM (15 minutes)

The Fabric for Frontier Experiments (FIFE) project is a major initiative within the Fermilab Scientific Computing Division charged with leading the computing model for Fermilab experiments. Work within the FIFE project creates close collaboration between experimenters and computing professionals to serve high-energy physics experiments of differing size, scope, and physics area. The FIFE project has worked to develop common tools for job submission, certificate management, software and reference data distribution through CVMFS repositories, robust data transfer, job monitoring, and databases for project tracking. Since the project's inception the experiments under the FIFE umbrella have significantly matured, and present an increasingly complex list of requirements to service providers. To meet these requirements, the FIFE project has been involved in transitioning the Fermilab General Purpose Grid cluster to support a partitionable slot model, expanding the resources available to experiments via the Open Science Grid, assisting with commissioning dedicated high-throughput computing resources for individual experiments, supporting the efforts of the HEP Cloud projects to provision a variety of back end resources, including public clouds and high performance computers, and developing rapid onboarding procedures for new experiments and collaborations. The larger demands also require enhanced job monitoring tools, which the project has developed using such tools as ElasticSearch and Grafana. FIFE has also closely worked with the Fermilab Scientific Computing Division's Offline Production Operations Support Group (OPOS) in helping experiments manage their large-scale production workflows. This group in turn requires a structured service to facilitate smooth management of experiment requests, which FIFE provides in the form of the Production Operations Management Service (POMS). POMS is designed to track and manage requests from the FIFE experiments to run particular workflows, and support troubleshooting and triage in case of problems. Recently we have started to work on a new certificate management infrastructure called Distributed Computing Access with Federated Identities (DCAFI) that will eliminate our dependence on a specific third-party Certificate Authority service and better accommodate FIFE collaborators without a Fermilab Kerberos account. DCAFI integrates the existing InCommon federated identity infrastructure, CILogon Basic CA, and a MyProxy service using a new general purpose open source tool. We will discuss the general FIFE onboarding strategy, progress in expanding FIFE experiments' presence on the Open Science Grid, new tools for job monitoring, the POMS service, and the DCAFI project. We will also discuss lessons learned from collaborating with the OPOS effort and how they can be applied to improve efficiency in current and future experiment's computational work.

Tertiary Keyword (Optional)

Distributed workload management

Secondary Keyword (Optional)

Computing middleware

Primary Keyword (Mandatory)

Computing models

Authors: MAZZACANE, Anna (FNAL); DYKSTRA, Dave (Fermi National Accelerator Lab. (US)); GARZOGGIO, Gabriele; Ms TEHERAN, Jeny (Fermilab); BOYD, Joseph (Fermilab); Dr HERNER, Kenneth Richard (Fermi National Accelerator Laboratory (US)); RETZKE, Kevin (Fermilab); MENGEL, Marc (Fermilab); VOTAVA, Margaret (Fermilab); KIRBY, Michael (Fermi National Accelerator Laboratory); ALTUNAY, Mine (Fermi National Accelerator Laboratory); SHARMA, Neha (Fermilab); LEVSHINA, Tanya

Presenter: Dr HERNER, Kenneth Richard (Fermi National Accelerator Laboratory (US))

Session Classification: Track 3: Distributed Computing

Track Classification: Track 3: Distributed Computing