Contribution ID: **394**   Type: **Talk**

Optimization of distributed compute resources utilization in the CMS Global Pool

Wednesday 23 October 2024 17:45 (18 minutes)

Efficient utilization of vast amounts of distributed compute resources is a key element in the success of the scientific programs of the LHC experiments. The CMS Submission Infrastructure is the main computing resource provisioning system for CMS workflows, including data processing, simulation and analysis. Resources geographically distributed across numerous institutions, including Grid, HPC and cloud providers, are joined into a set of federated resource pools, supervised by HTCondor and GlideinWMS services. The CMS Submission Infrastructure team is responsible for acquiring and managing this aggregated computing power, with a total capacity of about 500k CPU cores, and assigning it to CMS workloads according to their requirements and the priorities defined by the collaboration.

The scheduling strategies implemented for this purpose need to be flexible enough to support a number of concurrent workload types, taking into account the availability of resources from diverse providers, as well as the evolving resource requirements of the processing campaigns that the system needs to manage concurrently and consecutively. This complex system needs to be optimized in order to maximize the resource utilization efficiency, thus harnessing the full potential of our distributed compute resources.

This contribution will describe the systematic investigation by the CMS Submission Infrastructure team aimed at identifying, classifying, and minimizing inefficiencies in the use of the CMS distributed resources resulting from our workload management and scheduling algorithms. Additionally, our presentation will include certain strategies devised and implemented to compensate for other sources of inefficiency, thereby optimizing resource utilization and enhancing the overall CMS computational throughput.

Primary authors: PEREZ-CALERO YZQUIERDO, Antonio (Centro de Investigaciones Energéticas Medioambientales y Tecnológicas);   VON CUBE, Florian (KIT - Karlsruhe Institute of Technology (DE));   KIM, Hyunwoo (Fermi National Accelerator Lab. (US));   MASCHERONI, Marco (Univ. of California San Diego (US));   ZOKAITE, Vaiva (Vilnius University (LT))

Co-author: WURTHWEIN, Frank (UCSD)

Presenter: MASCHERONI, Marco (Univ. of California San Diego (US))

Session Classification: Parallel (Track 4)

Track Classification: Track 4 - Distributed Computing