



Contribution ID: 315

Type: Poster

Evolution of the CMS Submission Infrastructure to support heterogeneous resources in the LHC Run 3

Monday 24 October 2022 16:10 (30 minutes)

The landscape of computing power available for the CMS experiment is rapidly evolving, from a scenario dominated by x86 processors deployed at WLCG sites, towards a more diverse mixture of Grid, HPC, and Cloud facilities incorporating a higher fraction of non-CPU components, such as GPUs. Using these facilities' heterogeneous resources efficiently to process the vast amounts of data to be collected in the LHC Run3 and beyond, in the HL-LHC era, is key to CMS's achieving its scientific goals.

The CMS Submission Infrastructure is the main computing resource provisioning system for CMS workflows, including data processing, simulation and analysis. It currently aggregates nearly 400k CPU cores distributed worldwide from Grid, HPC and cloud providers. The Submission Infrastructure, together with other elements in the CMS workload management, has been modified in its strategies and enlarged in its scope to make use of these new resources.

In this evolution, key questions such as the optimal level of granularity in the description of the resources, or how to prioritize workflows in this new resource mix must be taken into consideration. In addition, access to many of these resources is considered opportunistic by CMS, thus each resource provider may also play a key role in defining particular allocation policies, diverse from the up-to-now dominant system of pledges. All these matters must be addressed in order to ensure the efficient allocation of resources and matchmaking to tasks to maximize their use by CMS.

This contribution will describe the evolution of the CMS Submission Infrastructure towards a full integration and support of heterogeneous resources according to CMS needs. In addition, a study of the pool of GPUs already available to CMS Offline Computing will be presented, including a survey of their diversity in relation to CMS workloads, and the scalability reach of the infrastructure to support them.

Significance

The Submission Infrastructure is the main component of the resource acquisition and workload to resource matchmaking systems in CMS Offline Computing. It is therefore mandatory to adapt it to be able to send GPU allocation requests to resource providers, to integrate those GPUs into the CMS HTCondor infrastructure, and finally to optimize workload to heterogeneous resource assignment in order for CMS to succeed in this future vast amount of computing power available in the form of GPUs. This contribution will present how this has been achieved, and indeed the already existing pool of GPUs ready for CMS use.

Experiment context, if any

The CMS experiment at the LHC at CERN

References

Primary author: PEREZ-CALERO YZQUIERDO, Antonio (Centro de Investigaciones Energéticas Medioambientales y Tecnológicas)

Co-authors: KIZINEVIC, Edita (CERN); KHAN, Farrukh Aftab (Fermi National Accelerator Lab. (US)); KIM, Hyunwoo (Fermi National Accelerator Lab. (US)); MASCHERONI, Marco (Univ. of California San Diego (US)); ACOSTA FLECHAS, Maria (Fermi National Accelerator Lab. (US)); HALEEM, Saqib (National Centre for Physics (PK))

Presenter: PEREZ-CALERO YZQUIERDO, Antonio (Centro de Investigaciones Energéticas Medioambientales y Tecnológicas)

Session Classification: Poster session with coffee break