

HEPCloud, an elastic virtual cluster from heterogeneous computing resources

Tuesday 26 April 2022 18:00 (25 minutes)

Historically, high energy physics computing has been performed on large purpose-built computing systems. These began as single-site compute facilities, but have evolved into the distributed computing grids used today.

The current computing landscape is more heterogeneous because of the elevated capacity and capability of commercial clouds and the push of funding agencies toward supercomputers. Both add new complications. Commercial cloud resources are highly virtualized and customizable but need to be managed. High Performance Computers are each one of a kind with different access rules and restrictions, like limited network connectivity or complex access patterns.

HEPCloud is a single managed portal that allows more scientists, experiments, and projects to use more resources to extract more science. Its goal is to provide cost-effective access by optimizing usage across all available types of computing resources and elastically expand the resource pool on short notice (e.g. by renting temporary resources on commercial clouds).

Fermilab HEPCloud facility has been used successfully in production for over three years providing and 2021 saw a big ramp up, especially for CMS that used all its Frontera quota 6 months ahead of expiry and used 90M NERSC-hours bonus after consuming all its allocation.

The Decision Engine is the software at the heart of HEPCloud, deciding where and how much to provision. It is an open-source project (<https://github.com/HEPCloud/decisionengine>) and recently version 2.0 was released, a release that we consider ready for wider adoption: it has a simplified installation and configuration, it is fully Python 3 code with strict coding best practices, it has a revised architecture with robust message passing between the components making the decisions.

Desired slot length

Speaker release

Yes

Authors: NORMAN, Andrew; MAMBELLI, Marco (Fermilab (US)); TIMM, Steven (Fermi National Accelerator Lab. (US))

Presenter: MAMBELLI, Marco (Fermilab (US))

Session Classification: Computing & Batch Services

Track Classification: Computing & Batch Services