

21st International Conference on Computing in High Energy and Nuclear Physics (CHEP2015)



Contribution ID: 112

Type: oral presentation

Jobs masonry with elastic Grid Jobs

Monday 13 April 2015 17:00 (15 minutes)

The DIRAC workload management system used by LHCb Distributed Computing is based on Computing Resource reservation and late binding (also known as pilot job in the case of batch resources) that allows the serial execution of several jobs obtained from a central task queue. CPU resources can usually be reserved for limited duration only (e.g. batch queue time limit) and in order to optimize their usage, it is important to be able to use them for the whole available time. However traditionally the tasks to be performed by jobs are defined at submission time and therefore it may happen that no job fits in the available time.

In LHCb, this so-called job masonry is optimized by the usage of elastic simulation jobs: unlike data processing jobs that must process all events in the input dataset, simulation jobs offer an interesting degree of freedom as one can define the number of events to be simulated even after submission. This requires however knowing three information: the time available in the reserved resource, its CPU power and the average CPU work required for simulating one event. The decision on the number of events can then be made at the very last moment, just before starting the simulation application. This is what we call elastic jobs.

LHCb simulation jobs are now all elastic, with an upper limit on the number of events per job. When several jobs are needed to complete a simulation request, enough jobs are submitted for simulating the total number of events required, assuming this upper limit for each job. New jobs are then submitted depending on the actual number of events simulated by this first batch of jobs.

We will show in this contribution how elastic jobs allow a better backfilling of computing resources as well as using resources with limited work capacity, such as short batch queues or volunteer computing resources. They also allow easily to shutdown virtual machines on cloud resources when sites require them to shutdown within a grace period.

Author: STAGNI, Federico (CERN)

Co-author: CHARPENTIER, Philippe (CERN)

Presenter: STAGNI, Federico (CERN)

Session Classification: Track 4 Session

Track Classification: Track4: Middleware, software development and tools, experiment frameworks, tools for distributed computing