

# Merging experimental and production clusters

Fabrice Huet  
`fabrice.huet@unice.fr`

30 janvier 2015

**Nombre d'étudiants souhaités : 3-4**

## Description du sujet

Researchers regularly use remote computing resources as part of their research. A lot of research organizations have in-house clusters which are easily accessible, with few limitations. Depending on the type of research, the number of nodes and the duration of the computation vary. For example, testing some distributed algorithm might require a large number of nodes for a few hours, whereas performing an astronomical simulation might last weeks. To deal with such a wide spectrum of requirements, clusters are usually set-up for a particular workload. A research cluster will allow jobs of short duration, but on a large number of nodes. On the other hand, a production cluster will allow very long running jobs on a few nodes to give every user a chance. These limitations are usually described in a charter and enforced by a scheduler. Very large Cloud providers like Amazon don't have such a problem because of the "infinite" number of nodes they provide and the associated cost for the user. However, a university cannot afford this.

The aim of this work is to answer the following question : is there a set of rules which would allow both experimental and production workloads on a single platform ?

To answer this question, we will have access to the logs of two job schedulers (experimental and production). The first part of the work will hence consist in data-mining these logs to extract relevant information (job duration, resources requested, resources granted...). In the second part, these information will be used to design an ideal cluster which allows for the execution of the most of the jobs, with a limited cost. This will be done mostly through simulation using the CloudSim framework

## Lieu

Université et INRIA.

## Prérequis

Bon niveau de Java. Pas de connaissance requise en Cluster ou Systèmes Distribués.

## Informations complémentaires