

_5vAW.png _5vAW.bb _5vAW.png”

Contribution ID: 132

Type: **Poster presentation**

Estimating job runtime for CMS analysis jobs

Monday, 14 October 2013 15:00 (45 minutes)

The basic premise of pilot systems is to create an overlay scheduling system on top of leased resources. And by definition, leases have a limited lifetime, so any job that is scheduled on such resources must finish before the lease is over, or it will be killed and all the computation wasted. In order to effectively schedule jobs to resources, the pilot system thus requires the expected lifetime of the jobs. Past studies have shown that relying on user provided estimates is not a valid strategy, so the system should try to make an estimate by itself. This paper provides a description of a system that makes estimates using machine learning based on past behavior. The work was performed in the context of physics analysis jobs of the CMS experiment at the Large Hadron Collider, using the domain knowledge to improve the accuracy. The attained results are presented in the paper.

Summary

Primary author: Mr SFILIGOI, Igor (University of California San Diego)

Presenter: Mr SFILIGOI, Igor (University of California San Diego)

Session Classification: Poster presentations

Track Classification: Distributed Processing and Data Handling A: Infrastructure, Sites, and Virtualization