



Contribution ID: 280

Type: **Poster presentation**

AutoPyFactory and the Cloud: Flexible, scalable, and automatic management of virtual resources for ATLAS

Monday, October 14, 2013 3:00 PM (45 minutes)

AutoPyFactory (APF) is a next-generation pilot submission framework that has been used as part of the ATLAS workload management system (PanDA) for two years. APF is reliable, scalable, and offers easy and flexible configuration. Using a plugin-based architecture, APF polls for information from configured information and batch systems (including grid sites), decides how many additional pilot jobs are needed, and submits them.

With the advent of cloud computing, providing resources goes beyond submitting pilots to grid sites. Now, the resources on which the pilot will run also need to be managed. Handling both pilot submission and controlling the virtual machine life cycle (creation, retirement, and termination) from the same framework allows robust and efficient management of the process.

In this paper we describe the design and implementation of these virtual machine management capabilities of APF. Expanding on our plugin-based approach, we allow cascades of virtual resources associated with a job queue. A single workflow can be directed first to a private, facility-based cloud, then a free academic cloud, then spot-priced EC2 resources, and finally on-demand commercial clouds. Limits, weighting, and priorities are supported, allowing free or less expensive resources to be used first, with costly resources only used when necessary. As demand drops, resources are drained and terminated in reverse order. Performance plots and time series will be included, showing how the implementation handles ramp-ups, ramp-downs, and spot terminations.

Primary author: HOVER, John (Brookhaven National Laboratory (BNL)-Unknown-Unknown)

Co-authors: Dr CABALLERO BEJAR, Jose (Brookhaven National Laboratory (US)); LOVE, Peter (LANCASTER UNIVERSITY)

Presenter: HOVER, John (Brookhaven National Laboratory (BNL)-Unknown-Unknown)

Session Classification: Poster presentations

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