

An Adaptive Batch Environment for Clouds

Wednesday, 21 April 2010 14:30 (30 minutes)

We have developed a method using Condor and a simple software component we call Cloud Scheduler to run High Throughput Computing workloads on multiple clouds situated at different sites. Cloud Scheduler is able to instantiate user-customized VMs to meet the requirements of jobs in Condor queues. Users supply Condor job attributes to indicate the type of VM required to complete their job. Condor then matches jobs with the appropriate VMs. Cloud Scheduler supports Nimbus and Amazon EC2, in addition we have preliminary support for Eucalyptus, and we plan support for OpenNebula 1.4. We present our early experiences with two separate projects using this method.

Primary author: GABLE, Ian (University of Victoria)

Presenter: GABLE, Ian (University of Victoria)

Session Classification: Virtualization