



Contribution ID: 94

Type: oral presentation

glideinWMS - A generic pilot-based Workload Management System

Tuesday, September 4, 2007 11:40 AM (20 minutes)

Grids are making it possible for Virtual Organizations (VOs) to run hundreds of thousands of jobs per day. However, the resources are distributed among hundreds of independent Grid sites. A higher level Workload Management System (WMS) is thus necessary.

glideinWMS is a pilot-based WMS, inheriting several useful features:

- 1) Late binding: Pilots are sent to all suitable Grid sites. Only once pilots start are real jobs selected for that resources. No forecasting is needed.
- 2) Reliability: A broken Grid site will either kill pilot jobs or pilots will detect the problem at startup. Real jobs only start on well-behaved resources.
- 3) Grid-wide fair share: The relative priorities between jobs of the same VO are set inside the WMS. Grid sites only manage priorities between different VOs.

glideinWMS is based on the Condor glidein concept, i.e. a regular Condor pool, with the Condor daemons (startd) being started by pilot jobs. The real jobs are vanilla, standard or MPI universe jobs.

glideinWMS is composed of Glidein Factories and VO Frontends, communicating using Condor ClassAds:

- * Factories publish the available Grid sites,
- * Frontends match the Grid attributes to job attributes and publish a request for a stream of glideins to suitable Grid sites
- * Factories pick up the requests and submit the glideins

A detailed description of the system will be presented, along with the currently deployed systems inside USCMS production and user analysis frameworks. Integration with frameworks of other VOs will also be presented, as well as the measured scalability limits.

Author: Mr SFILIGOI, Igor (FNAL)

Presenter: Mr SFILIGOI, Igor (FNAL)

Session Classification: Grid middleware and tools

Track Classification: Grid middleware and tools