

CorralWMS: Integrated Resource Provisioning Across the National Cyberinfrastructure in Support of Scientific Workloads

Ewa Deelman, USC; Burt Holzman, Fermi Laboratory; Frank Wuerthwein, USCD

The goal of this work is to develop and deploy integrated, robust, and scalable resource provisioning services that support a broad set of domain application workflow and workload execution environments. The aim is to integrate and enable these services across local and distributed computing resources, the major national cyberinfrastructure providers, Open Science Grid and TeraGrid, as well as emerging commercial and community cloud environments.

CorralWMS extends and reuses existing solutions that are already proven in production. Both Corral and glideinWMS currently operate as standalone resource provisioning systems. GlideinWMS was initially developed to meet the needs of the CMS (Compact Muon Solenoid) experiment at the LHC at CERN. It generalizes a Condor GlideIn system developed for CDF (The Collider Detector at Fermilab) and first deployed for production in 2003. It has been in production across the Worldwide LHC Computing Grid (WLCG), with major contributions from the Open Science Grid. GlideinWMS has been used in production with 8,000 concurrently running jobs, totaling more than six million jobs executed over the last year, consuming more than 2000 CPU years.

Corral, a static resource provisioning system was recently built to meet the needs of workflow-based applications running on the TeraGrid. It is being used today by the Southern California Earthquake Center (SCEC) CyberShake application. In a period of 10 days in May 2009, SCEC used Corral to provision a total of 33,600 cores and used them to execute 50 workflows, each containing approximately 800,000 application tasks, which corresponded to 852,120 individual jobs executed on the TeraGrid Ranger system.

The proposed extended integrated system, **CorralWMS**, will retain the identity of both systems to maintain backward compatibility for existing applications that use either Corral or glideinWMS, while new applications will benefit from the expanded set of provisioning capabilities of both systems combined. As with the current solutions, job management on the provisioned resources will be done via Condor, which allows for the customization of job scheduling policies.

This new system combines the benefits of the proactive provisioning performed by Corral with the dynamic resource provisioning performed by glideinWMS. In addition, the proposed program of work will extend the existing provisioning options to include cloud computing resources, which are of interest to the current stakeholder communities. Finally, the monitoring capabilities will be improved—extending the ability to monitor overall system behavior (the provisioning being performed) as well as the jobs making use of these resources.

