

Recent developments in glideinWMS: minimizing resource wastages

Tuesday, July 10, 2018 4:40 PM (20 minutes)

GlideinWMS is a workload management system that allows different scientific communities, or Virtual Organizations (VO), to share computing resources distributed over independent sites. A dynamically sized pool of resources is created by different VO-independent glideinWMS pilot factories, based on the requests made by the several VO-dependant glideinWMS frontends. For example, the CMS VO employs a pool that ranges from 80k to over 200k cores, worldwide distributed over 50 sites. With such big numbers, calculating the proper pilot pressure is essential to efficiently provision resources, and being able to effectively drain resources during site downtimes also helps minimize the wastages in the system. However, the often spiky nature of the demand, and the geometry of certain sites that may have multiple entry points, makes it difficult to tightly couple pilot submission to the actual demand, especially in a multicore environment. In this contribution we discuss recent enhancements to glideinWMS that aim at improving the efficiency of the resource provisioning by adding the concept of metasite, and by introducing new handles the VOs can use to manage the pilot pressure. New ways of draining sites using the Machine/Job Feature mechanism and plans for the future are also discussed.

Primary authors: MASCHERONI, Marco (Univ. of California San Diego (US)); LETTS, James (Univ. of California San Diego (US)); PEREZ-CALERO YZQUIERDO, Antonio (Centro de Investigaciones Energéticas Medioambientales y Tecnológicas); MHASHILKAR, Parag (Fermi National Accelerator Laboratory); MAMBELLI, Marco (University of Chicago (US)); DOST, Jeffrey Michael (Univ. of California San Diego (US)); BOX, Dennis (F)

Presenter: MASCHERONI, Marco (Univ. of California San Diego (US))

Session Classification: Posters

Track Classification: Track 3 –Distributed computing