

Pilot Factory - a Condor-based System for Scalable Pilot Job Generation in the Panda WMS Framework

Thursday, March 26, 2009 8:00 AM (20 minutes)

The Panda Workload Management System is designed around the concept of the Pilot Job - a “smart wrapper” for the payload executable, that can probe the environment on the remote worker node before pulling down the payload from the server and executing it. Such design allows for improved logging and monitoring capabilities as well as flexibility in Workload Management. In the Grid environment (such as the Open Science Grid), Panda Pilot Jobs are submitted to remote sites via mechanisms that ultimately rely on Condor-G. As our experience has shown, in cases where a large number of Panda jobs are simultaneously routed to a particular remote site, the increased load on the head node of the cluster, which is caused by the Pilot Job submission, may lead to overall lack of scalability. We have developed a Condor-inspired solution to this problem, which is using the schedd-based glidein, whose mission is to redirect pilots to the native batch system. Once a glidein schedd is installed and running, it can be utilized exactly the same way as local schedds and therefore, from the user’s perspective, Pilots thus submitted are quite similar to jobs submitted to the local Condor pool.

Presentation type (oral | poster)

poster

Primary authors: CHIU, Barnett (BROOKHAVEN NATIONAL LABORATORY); Dr POTEKHIN, Maxim (BROOKHAVEN NATIONAL LABORATORY)

Presenter: Dr POTEKHIN, Maxim (BROOKHAVEN NATIONAL LABORATORY)

Session Classification: Poster session

Track Classification: Grid Middleware and Networking Technologies