Contribution ID: **35** Type: **oral presentation**

The Capone Workflow Manager

Tuesday, 14 February 2006 14:20 (20 minutes)

We describe the Capone workflow manager which was designed to work for Grid3 and the Open Science Grid. It has been used extensively to run ATLAS managed and user production jobs during the past year but has undergone major redesigns to improve reliablility and scalability as a result of lessons learned (cite Prod paper). This paper introduces the main features of the new system covering job management, monitoring, troublehsooting, debugging and job logging. Next, the modular architecture which implements several key evolutionary changes to the system is described: a multi-threaded pool structure, checkpointing mechanisms, and robust interactions with external components, all developed to address scalability and state persistence issues uncovered during operations running of the production system. Finally, we describe the process of delivering production ready tools, provide results from benchmark stress tests, and compare Capone with other workflow managers in use for distributed production systems.

Primary author: MAMBELLI, Marco (UNIVERSITY OF CHICAGO)

Co-authors: GIERALTOWSKI, Jerry (ANL (ARGONNE NATIONAL LABORATORY)); GARDNER, Robert (UNIVERSITY OF CHICAGO)

Presenter: MAMBELLI, Marco (UNIVERSITY OF CHICAGO)

Session Classification: Software Tools and Information Systems

Track Classification: Software Tools and Information Systems