



Contribution ID: 321

Type: oral presentation

SPHINX: A Scheduling Middleware for Data Intensive Applications on a Grid

Thursday, 30 September 2004 14:40 (20 minutes)

A grid consists of high-end computational, storage, and network resources that, while known a priori, are dynamic with respect to activity and availability. Efficient co-scheduling of requests to use grid resources must adapt to this dynamic environment while meeting administrative policies. We discuss the necessary requirements of such a scheduler and introduce a distributed framework called SPHINX that schedules complex, data intensive High Energy Physics and Data Mining applications in a grid environment, respecting local and global policies along with a specified level of quality of service. The SPHINX design allows for a number of functional modules and/or distributed services to flexibly schedule workflows representing multiple applications on grids. We present experimental results for SPHINX that effectively utilize existing grid middleware such as monitoring and workflow management/execution systems. These results demonstrate that SPHINX can successfully schedule work across a large number of grid sites that are owned by multiple units in a virtual organization.

Primary authors: IN, J.U. (UNIVERSITY OF FLORIDA); CHITNIS, L. (University of Florida); KULKARNI, M. (UNIVERSITY OF FLORIDA); AVERY, P. (UNIVERSITY OF FLORIDA); CAVANAUGH, R. (UNIVERSITY OF FLORIDA); RANKA, S. (UNIVERSITY OF FLORIDA)

Presenter: CAVANAUGH, R. (UNIVERSITY OF FLORIDA)

Session Classification: Distributed Computing Services

Track Classification: Track 4 - Distributed Computing Services