

Pseudo-interactive monitoring in distributed computing

Monday 23 March 2009 08:00 (20 minutes)

Distributed computing, and in particular Grid computing, enables physicists to use thousands of CPU days worth of computing every day, by submitting thousands of compute jobs.

Unfortunately, a small fraction of such jobs regularly fail; the reasons vary from disk and network problems to bugs in the user code. A subset of these failures result in jobs being stuck for long periods of time. In order to debug such failures, interactive monitoring is highly desirable; users need to browse through the job log files and check the status of the running processes.

Batch systems typically don't provide such services; at best, users get job logs at job termination, and even this may not be possible if the job is stuck in an infinite loop.

In this paper we present a novel approach of using regular batch system capabilities of Condor to enable users to access the logs and processes of any running job. This does not provide true interactive access, so commands like vi are not viable, but it does allow operations like ls,

cat, top, ps, lsof, netstat and dumping the stack of any process owned by the user; we call this pseudo-interactive monitoring.

It is worth noting that the same method can be used to monitor Grid jobs in a glidein-based environment.

We further believe that the same mechanism could be applied to many other batch systems.

Author: Mr SFILIGOI, Igor (Fermilab)

Co-authors: Mr BRADLEY, Daniel (University of Wisconsin, Madison); Prof. LIVNY, Miron (University of Wisconsin, Madison)

Presenter: Mr SFILIGOI, Igor (Fermilab)

Session Classification: Poster session

Track Classification: Distributed Processing and Analysis