

7B4 The Grid Job Monitoring Service

Ludek Matyska, Ales Krenek, Miroslav Ruda, Michal Vocu, Zdenek Salvat, Jiri Sitera, Jan Pospisil and Daniel Kouril, CESNET, Czech Republic

Abstract:

The Logging and Bookkeeping (LB) service is currently being developed within the DataGRD Project to solve the problem of the job monitoring within the heterogeneous Grid environment, where users lose control over job submission and are also unable, for security reasons, to directly contact the computing elements where jobs are run.

The LB service collects important events in the job life (e.g., a matching resource found, job execution started, etc.), delivers them to a LB server and stores them in a reliable way. In general, an event indicates a change in job status. The LB service architecture features two APIs, local logger subservice and the servers. Events generated by the Grid components enter the LB system through the logging (producer) API. This API, implemented as a simple library, provides locally persistent non-blocking calls, regardless of the state of other LB components. A message describing the event is passed to the local logger subservice. Local persistence is implemented with a transaction log file written by the locallogger, while the interlogger is responsible for actual transfer of messages to the LB databases, typically over WAN. The LB servers manage the persistent data storage. The event is the principal storage entity, job status is computed on demand, by analysing the event, when client query the server with a server (producer) API. Jobs are identified by a unique Grid job identifier, whose syntax allows the retrieval of virtually any information on the job by direct access to the bookkeeping server, without the need to contact another Grid service. In fact, the bookkeeping server speaks a subset of the HTTPS protocol; information on job status is obtained with its GET request.

All the communication between LB components is authenticated with either a user or a service certificate. All the queries are also authenticated and currently only information about user's own jobs is provided. The LB system just described is a part of the resource management subsystem for DataGrid project currently used and evaluated on the testbed. Experience with the first version led to extensions with notification, user attributes and dynamic and multiple job support.

.