



Contribution ID: 89

Type: **Oral Presentation**

Logging and Bookkeeping –What Can it Do for You?

Monday, 11 April 2011 16:30 (30 minutes)

Overview

gLite's Logging and Bookkeeping service (L&B) has been with us for almost 10 years as a robust monitoring service which gathers, aggregates and archives information on infrastructure behavior from the perspective of users' tasks, and provides a single (per VO) endpoint for accessing that information. Numerous use cases have emerged over the years, some well-known and widely applied in gLite-based grids, some not so widely known but possibly important for the user community.

The current situation gives us an opportunity to stop and take time to review L&B's potential for various usage scenarios and to plan its continued development according to real-world needs. We would like to involve current and potential L&B users (end users, groups, VO managers, monitoring tool developers/designers) in this process. To help this happen we give an outline of current L&B features and use cases (sometimes not widely known but ready to use or prototyped) and possible directions of further development.

Impact

The main purpose of this talk is to trigger discussion with the user community. It gives an overview of existing usage patterns involving the L&B service, emphasizing less-well known use cases and applications often well outside the scope of computing job tracking. We invite users and user communities to provide their input. They are welcome to approach the L&B team with their priorities and requests to support other interesting use cases.

As a starting point for the discussion we present several topics L&B wishes to address in the near future. Feedback on priorities or possible extensions of the list is more than welcome.

- interesting entities to be tracked, besides gLite WMS jobs (e.g. ARC and UNICORE CE jobs, data transfers, SRM operations, ...),
- the ability to track dependencies among such entities (e.g. a computational job is blocked by waiting transfer of its input),
- the desired level of complexity of queries (all jobs of a particular user, the user's jobs within a given time interval, failed/successful jobs on a given CE, ..., up to full SQL/XQuery/etc. power on the job data)
- output formats to be supported (Glue-conformant WS interface, simple key=value text format, human-readable HTML, ...),
- history of the data to be kept (a day, a week, a month, ...),
- types of aggregate information, e.g. average queue traversal time, job failure rate etc., level of aggregation (per user, VO, grid service instance).

Description of the work

The notoriously known usage of L&B are the `glite-wms-job-status` and `glite-wms-job-logging-info` commands, querying status of jobs handled by gLite WMS. It is less known that monitoring applications such as RTM, Experiment Dashboard or Grid observatory heavily rely on L&B, too.

Starting with its recent release 2.1, L&B is able to track native (i.e. non-WMS) CREAM jobs. Experimental support for PBS jobs and Condor is also available. L&B can also work with data transfers (e.g. sandbox transfers), and the job status information is enhanced with the state of associated data movements. We believe that a great opportunity to provide L&B users with a more accurate and descriptive view of their jobs lies in tracking the dependencies between computational and data transfer jobs, either directly through job state L&B queries or indirectly through monitoring tools providing VO or site specific views.

L&B can store application-specific info (job annotations, metrics, or status) as user tags. The tags can be queried, and used to build application-specific dashboards, as we demonstrated several times. We also provide support for tracking generic user workflows currently used in medical image processing (a subject of a standalone contribution).

We can also demonstrate successful use of L&B to track different entities –CA revocation lists. Individual CRLs are registered in L&B in place of jobs. Whenever the CRL is updated, it is reflected in state update of the corresponding “job” in L&B, which triggers delivery of L&B notifications to the sites subscribed for receiving CRL updates. The payload of the notification carries the actual CRL update then.

URL

<http://egee.cesnet.cz/cs/JRA1/LB/>

Conclusions

The grid environment matured in the recent years and essential functionality of the services is, more or less, available. The current development activities, besides pursuing standardization and harmonization of parallel solutions, aim at providing advanced functionality, which can be leveraged by the emerging grid user communities.

In this talk we present, besides a summary of less-known usage patterns of the service, a view on possible directions of further development of the L&B service as we ourselves can foresee its advanced use. The talk is intended to trigger discussion with the user community, which will result into a more specific and eventually extended work plan, better aligned with the expected user needs.

Primary authors: Dr KRENEK, Ales (CESNET); Mr KOURIL, Daniel (CESNET); Mr DVORAK, Frantisek (CESNET); Mr FILIPOVIC, Jiri (CESNET); Mr SITERA, Jiri (CESNET); Prof. MATYSKA, Ludek (CESNET); Mr SUSTR, Zdenek (CESNET)

Presenters: Mr SITERA, Jiri (CESNET); Mr SUSTR, Zdenek (CESNET)

Session Classification: User Support Services

Track Classification: User Support Services - Application/Community