

Logging and Bookkeeping

What Can it Do for You

Zdeněk Šustr et al.

CESNET

11 April 2011



Agenda

- ▶ Introduction to the Logging and Bookkeeping service (**L&B**)
- ▶ Recent developments
 - ▶ Final release with EGEE (2.1)
 - ▶ EMI-1 *Kebnekaise* release (3.0)
- ▶ Further development in the future, near and distant

L&B – an Introduction

- ▶ Started life in EDG/EGEE
 - ▶ part of the gLite middleware stack
- ▶ A service for job-centric monitoring on the grid
 - ▶ collecting events from grid elements
 - ▶ using that information to reconstruct job state
 - ▶ presenting a single point of contact for users to...
 - ▶ query the state of their jobs
 - ▶ register for automated notifications on job state changes
 - ▶ register user tags
- ▶ Reliable delivery chain for asynchronous delivery of events

Job States

- ▶ State machine designed to recover from missing events, resubmissions. . .
- ▶ Multiple job types supported
 - ▶ Native gLite WMS jobs (depicted)
 - ▶ Sandbox transfers
 - ▶ PBS jobs
 - ▶ Direct submissions to CREAM
 - ▶ Condor
- ▶ Even more job types planned
 - ▶ other EMI middleware stacks
 - ▶ generic state diagram

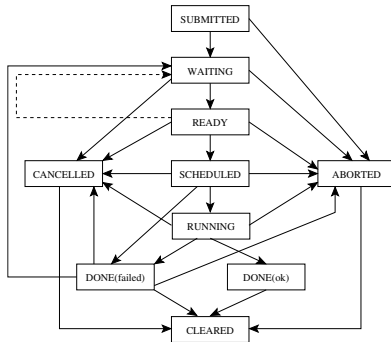


Fig. 1: gLite WMS job state diagram

User Interaction

- ▶ **query/response** – standard user queries with L&B's client API
 - ▶ Limited capabilities also accessible over *HTTP*s or *WS*
- ▶ **publish/subscribe** – avoiding repeated queries: register once, specify conditions, receive notification messages
 - ▶ Accessible either through L&B's client API or
 - ▶ Recently developed capability to deliver to **ActiveMQ**
 - ▶ See Poster No. 5 (L&B and Messaging) for details

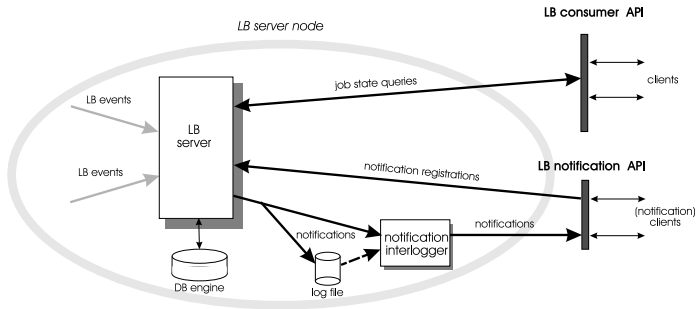


Fig. 2: L&B user access models

Recent Additions

- ▶ L&B 2.1 – Final product of EGEEIII
 - ▶ Adoption of the common logging format
 - ▶ Advanced authorization with unified configuration
 - ▶ Elementary native support for CREAM jobs
 - ▶ Ability to log sandbox transfer progress as a specific job type
 - ▶ Per-CE statistics used by WMS

- ▶ L&B 3.0 – Part of EMI-1 Kebnekaise: release imminent
 - ▶ Extended authorization features
 - ▶ Controlling permissions to log user tags
 - ▶ Support for changing job owners
 - ▶ Publishing L&B notification messages through **ActiveMQ**
 - ▶ Collections of sandbox transfers

Use Cases

- ▶ Monitoring grid jobs, collections, workflows
 - ▶ Widely deployed across the infrastructure (former EGEE grid)
- ▶ Producing higher-level information for dashboards
 - ▶ RTM (Real-Time Monitor), Imperial College London
 - ▶ LHC Experiment Dashboard, CERN
 - ▶ Grid Observatory, LAL, ASR, DiGITEO, Paris-sud-11
- ▶ Monitoring other automated workflows outside the scope of grid computing
 - ▶ Pilot project runs jointly with the Medical Faculty of the Masaryk University in Brno, producing the *Hypertext Atlases of Pathology*, Josef Feit et al., www.atlases.muni.cz
- ▶ Monitoring installation/update status across distributed systems
 - ▶ Monitoring status of security updates across a large infrastructure.
 - ▶ Monitoring the status of distribution of certificate revocation lists across grid infrastructure; MetaCentrum, the Czech NGI.
 - ▶ Collection of SysLog records from the national grid; MetaCentrum, the Czech NGI.

Near Future

- ▶ Receiving events from **ActiveMQ**
 - ▶ Allowing other grid elements/middleware stacks to log events to L&B without having to incorporate the L&B API.
- ▶ Tunneling **ActiveMQ** messages through L&B's delivery chain
 - ▶ Taking advantage of the fact that L&B clients are installed on **gLite**-enabled CEs
 - ▶ Allowing user jobs to send out messages targeted to **ActiveMQ** in the same way user tags are logged
- ▶ Extended support for native CREAM jobs
- ▶ Common job submission and management methods across EMI
- ▶ Consolidation and harmonization of compute area clients/APIs

What Else to Address. . .

- ▶ Other entities to track besides gLite WMS/CREAM jobs
 - ▶ ARC and Unicore jobs, data transfers, SRM operations. . .
- ▶ Ability to track dependencies between entities
 - ▶ e.g. a computational job is blocked by waiting for an input file transfer to finish
- ▶ Increased complexity of queries
 - ▶ extended querying capabilities, allowing for the use of (almost) SQL-like query conditions
- ▶ Extended support for additional output formats
 - ▶ WS (Glue), key=value format, human-readable HTML. . .
- ▶ History of the data to be kept
 - ▶ Days? Weeks? Months?
- ▶ Aggregate information
 - ▶ average queue traversal time, job failure rate etc.
 - ▶ level of aggregation (per user, VO, grid service instance).

Suggest, Prioritize

- ▶ There was a recent survey among EGI's SA1
- ▶ You are invited to add your suggestions and help us prioritize
- ▶ Interested in running a **joint effort to prototype your solution?** You are welcome to **contact us!**
- ▶ *emi-lb@metacentrum.cz*
- ▶ *<http://egee.cesnet.cz/en/JRA1/LB>*

The End

Thank you for your attention.