

Monitoring National Infrastructure with L&B

Description of the work

Besides its contribution to EGI, the Czech NGI operates its own distributed environment – the METACentrum. In addition to traditional computing clusters METACentrum has been providing virtualization services, either relying on its own virtualization infrastructure that allows creation of virtual clusters or extension of physical clusters with virtual nodes, or providing pure Open Nebula-based cloud services. Naturally the batch system can be used to schedule jobs for physical and virtual nodes alike. There are multiple instances of the batch system capable of forwarding jobs among themselves to improve load balancing or find alternatives in match making. L&B has been selected as a monitoring tool for the infrastructure for its ability to monitor processes whose events arrive from different nodes in a grid, to provide reliable message delivery and up-to-date status information regardless of occasional grid component or communication failures.

To achieve the goal L&B has been extended with state diagrams for torque jobs and virtual machines. Instrumentation of L&B logging calls from Torque has been implemented at the source code level, while virtualization stacks (both Open Nebula and home-brewed Magrathea) make use of existing callback hooks to log events through command line executables. Similar logging calls are made from Dom0 and also from within the virtual image, so that events triggering high-level state changes are redundant and distinguishing between similar events received from various sources provides for fine-grained status monitoring.

Another reason to choose a single monitoring tool for different processes is its ability to keep track of relationships. ID of a VM used to run a given computing job can be stored with the job. Similarly, with just a simple abstraction, VM state diagram can also be used for physical machines, and then a similar relationship can be established between a VM and the physical resources used to execute it.

Link for further information

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

Wider impact of this work

The work described here is intended not only as a production-level solution for METACentrum, but also as a prototype of a general-purpose cloud monitoring solution. Since L&B simply connects up to the infrastructure through messaging (supporting not only its own messaging protocol but also STOMP or OpenWire), it can be used to monitor any virtualized infrastructure capable of sending event messages over any of these channels. L&B also brings in other advanced features such as notifications or user tags that can be stored with other status information. It supports multiple authentication mechanisms, namely PKI (used across the EGI grid) and Kerberos (used in METACentrum).

The presented solution is also a prototype of integrating L&B's HTTPs interface into a Web portal to bring user-related information directly to the users' personal pages.

Printable Summary

gLite's Logging and Bookkeeping (L&B) is a monitoring tool equipped for monitoring the states of all kinds of processes related to grid and cloud computing. Besides traditional gLite WMS jobs and logical groupings thereof such as DAGs or collections it also supports input/output sandbox transfers, native CREAM jobs, Torque jobs and, as a recent addition, Virtual Machine states. With that, L&B has been deployed over the Czech NGI's infrastructure to provide users with a uniform view of all their processes, be it traditional jobs submitted to Torque-managed clusters, virtual machines managed through Czech NGI's own virtualization solution, or Open Nebula Cloud Machines. Where applicable, mutual links between jobs, virtual machines and physical machines are also recorded and made available as a part of the status information.

Primary authors: KOURIL, Daniel (CESNET); Mr DVORAK, Frantisek (CESNET); Mrs KRENKOVA, Ivana (CESNET); CHUDOBA, Jiri (CESNET); Mr FILIPOVIC, Jiri (CESNET); SITERA, Jiri (CESNET); MATYSKA, Ludek (CESNET); Mr POUL, Marcel (CESNET); VOCU, Michal; Mr RUDA, Miroslav (Cesnet); Mr TOTH, Simon (CESNET); SUSTR, Zdenek (CESNET)

Presenter: SUSTR, Zdenek (CESNET)

Track Classification: Resource Infrastructure services (Peter Solagna: track leader)