

Monitoring Virtual Machines with L&B – the FedCloud Experience

Tuesday, 9 April 2013 14:40 (25 minutes)

Impact

Notification streams produced by L&B in this manner are suitable for use by monitoring tools, dashboards, or community-specific workload distribution frameworks. Far from just translating events received over diverse input channels into a common format, L&B processes the information, employs its state diagram implementation to determine the current status and attributes of the given process regardless of possible irregularities in the input data such as events lost or delivered out of order, and then delivers the processed information over a selected channel to the target consumer.

Another L&B feature aimed specifically at supporting user community workflows allows for linking computing jobs, i.e., payload executing in the given virtual machine, to that machine's record in L&B, producing a comprehensive view of the computing process and underlying infrastructure. Thanks to this feature, L&B can track overall physical resource usage across the whole infrastructure regardless of end-user service type (cloud, grid job, virtual worker nodes). We believe that this will help infrastructure providers to really understand existing usage patterns in new virtualized environments.

URL

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

Summary

The Logging and Bookkeeping service (L&B), traditionally used to monitor grid computing jobs, has been recently extended to support also monitoring of virtual machine status. Aiming at capitalizing on the wide range of communication channels supported by L&B, it has been invited to provide cloud monitoring and notification services within EGI's FedCloud task. The FedCloud environment comprises of different virtualization stack implementations and deployment scenarios, requiring a common tool to monitor status and propagate standardized notifications in case of certain occurrences within the environment. Experimental use of L&B for that purpose started on national level, supporting only a limited subset of available solutions, and was already publicized. This talk details the experience made in extending the support to additional virtualization stacks and information sources.

Description

Similar to a traditional grid setup, where L&B combines information gathered from different grid elements into a single view of the current status of a computing job, individual components in a virtualization stack can be made sources of virtual-machine specific information (events). The most common event sources are the cloud or virtual cluster managers (OpenNebula, OpenStack, WNoDeS, Magrathea), virtualization hypervisors (Xen, KVM) and the virtual machine instances themselves. They all need to be instrumented to generate the appropriate events at appropriate times and deliver them to L&B. Alternatively, some of them already send out messages over one or more channels (one instance of this is OpenStack, which already publishes notifications over PubSubHubub) and L&B can tap into that stream to collect event messages.

L&B puts the collected information to several different uses. First of all, it stores it raw for a predefined period of time. Second, it processes it to determine the current state of the process (VM) in question, and its various attributes, and makes the information available over the querying interface (L&B API, HTTPs, ...). Third, if pre-configured conditions are met, it sends out the processed (!) information over a common notification channel in a unified format, regardless of the type, implementation or version of the original source.

The outgoing notification stream is highly configurable, allowing –on one hand –for notifications to be generated on any small event, or –on the other hand –only when a very complex set of conditions concerning the status and attributes of the given process is met.

Primary authors: PARAK, Boris (Masaryk University); DVORAK, Frantisek (CESNET); SITERA, Jiri (CESNET); FILIPOVIČ, Jiří (CESNET); VOCUS, Michal; Mr RUDA, Miroslav (Cesnet); SUSTR, Zdenek (CESNET)

Presenter: SUSTR, Zdenek (CESNET)

Session Classification: Cloud Platforms

Track Classification: Cloud Platforms (Track Lead: M Drescher and M Turilli)