

Monitoring Virtual Machine Status with L&B

*Šustr, Z., Sitera, J., Filipovič, J., Dvořák, F.,
Voců, M., Kouřil, D., Poul, M., Matyska, L.
CESNET, Czech Republic*

The **Logging and Bookkeeping** service

- Long-standing part of the gLite middleware (EDG, EGEE, EMI)
 - Monitoring production compute jobs, mostly WMS
 - Back-end to the well known `glite-wms-job-status` command
- Collects event messages from grid elements
- **Processes** the information
- Publishes status information in several ways:
 - *query/response* interface – L&B querying API/HTML
 - *subscribe/publish* interface – L&B notifications
 - ▶ OpenWire/STOMP, RSS, L&B legacy
- Supports multiple types of grid processes
 - Compute jobs, collections, file transfers
 - Newest addition: **Virtual Machines**

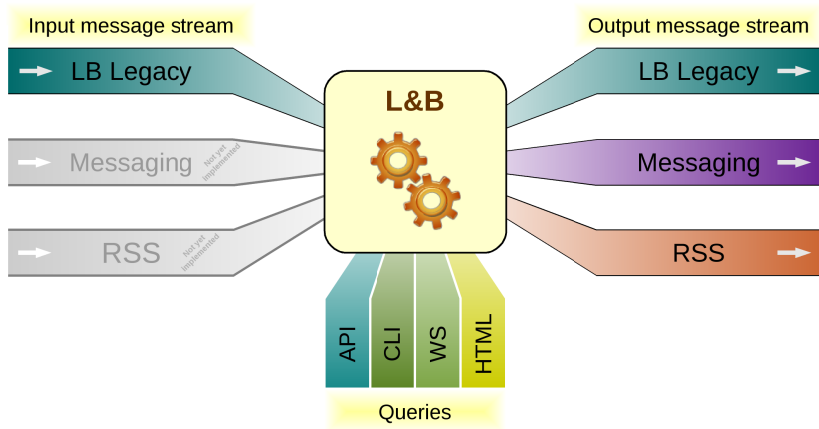


Fig. 1: L&B receiving and producing messages over different channels

Virtual Machine (VM) Life Cycle

- Obvious similarities with a compute job life cycle
 - Well-defined life span
 - State transitions triggered by event messages
 - Shared infrastructure
- Even more similarity with pilot jobs
 - Start up with no specific payload defined
 - Attract attention when ready, get assigned payload
- Added benefits of using grid monitoring tools
 - Already in place
 - The same tool (L&B) can be used to monitor both layers
 - Easy to keep track of VM-Payload relationship
- VM seen as **“just another job type”**

- Inspired by Open Nebula 3.0 specification
- Driven by events collected from multiple sources:
 - Open Nebula or another VM manager
 - Hypervisor such as *Xen*
 - Actual VM instance
- Event logging instrumented through
 - a client library (C, C++, Java)
 - command line tools
- VM-specific states mapped to generic gLite states for universal queries

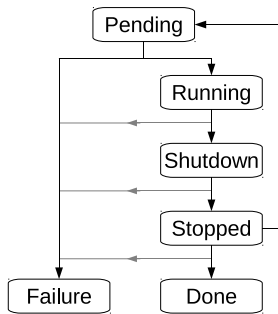


Fig. 2: VM state diagram

Component	Events	Attributes
CloudManager	Register, Create, Running, Host detail, Shutdown, Done	Hostname, Physical host name, Owner, Requirements...
VM Image	Really Running, Shutdown	Runtime information, user tags
Hypervisor	Running, Shutdown	

- Some types of events received from multiple components
 - Provide for redundancy and fine-grained monitoring
 - Fill in details known only at certain level

- Extension of known L&B concept
- Previously: specific case implementations
 - subjobs \leftrightarrow collection/DAG parents
 - compute jobs \leftrightarrow sandbox transfers
- New: **generic bilateral relationship**
 - Registered/updated by events
 - Relationship status record maintained
 - ▶ Distinguishes *active* from *past* relationships
 - One of intended uses: linking compute jobs to VMs
 - ▶ Typically registered by batch system, known to both jobs

Related Jobs

- <https://metalb.ics.muni.cz:9000/139.metalb.ics.muni.cz> – Pbs Relationship Active
- <https://metalb.ics.muni.cz:9000/107.metalb.ics.muni.cz> – Pbs Relationship Cancelled
- <https://metalb.ics.muni.cz:9000/140.metalb.ics.muni.cz> – Pbs Relationship Cancelled
- <https://metalb.ics.muni.cz:9000/142.metalb.ics.muni.cz> – Pbs Relationship Cancelled
- <https://metalb.ics.muni.cz:9000/143.metalb.ics.muni.cz> – Pbs Relationship Cancelled
- <https://metalb.ics.muni.cz:9000/144.metalb.ics.muni.cz> – Pbs Relationship Cancelled
- <https://metalb.ics.muni.cz:9000/145.metalb.ics.muni.cz> – Pbs Relationship Cancelled
- <https://metalb.ics.muni.cz:9000/146.metalb.ics.muni.cz> – Pbs Relationship Cancelled
- <https://metalb.ics.muni.cz:9000/147.metalb.ics.muni.cz> – Pbs Relationship Cancelled
- <https://metalb.ics.muni.cz:9000/148.metalb.ics.muni.cz> – Pbs Relationship Cancelled
- <https://metalb.ics.muni.cz:9000/149.metalb.ics.muni.cz> – Pbs Relationship Cancelled
- <https://metalb.ics.muni.cz:9000/150.metalb.ics.muni.cz> – Pbs Relationship Cancelled
- <https://metalb.ics.muni.cz:9000/151.metalb.ics.muni.cz> – Pbs Relationship Cancelled
- <https://metalb.ics.muni.cz:9000/152.metalb.ics.muni.cz> – Pbs Relationship Cancelled

Fig. 3: Jobs related to a VM – a screenshot

- Per-job ACL-based authorization
- Robust state interpretation
 - Resistant to events delivered out of order or not delivered at all
- Highly configurable output
 - Applies to synchronous queries as well as generated message streams
 - Complex sets of conditions supported
 - ▶ Messages sent over certain channels only if conditions are met

- Existing grid component has been extended to follow VMs as “another type of job”
- Keeping track of relationships with payload – regular compute jobs of different kinds
- Multiple channels available for receiving output, either by query or broadcast
- Potential for supporting event input over many channels

Thanks



Thank You