

AliEn Jobs in EC2 Computing Clouds

Maxim Storetvedt

Bergen University College
University of Bergen

31 March 2016



HØGSKOLEN
I BERGEN



Contents

1 Background

2 Project

3 Related Work

4 Questions

Background

Suggestion – Run AliEn jobs within a Cloud environment

- Platform can be provided by Openstack/OpenNebula/Nimbus etc.
- Support for commercial Clouds (Amazon EC2).
- Launch virtual worker nodes to execute jobs.

Benefits for ALICE

- Enable access to external resources.
- Ease the process of integrating new nodes.
- Provide elastic resource management.
 - Scale resources based on the number of jobs present.

Project

- Builds on an earlier project – AliEC2 by Joachim Carlsen¹.
- Framework for executing AliEn jobs in EC2 Clouds.
- Security mechanisms.
- Reducing VM startup overhead.

Why EC2 ?

- De-facto API interface for computing Clouds.
- Wide support (OpenStack, OpenNebula, Nimbus, Eucalyptus, etc.)

1. <http://github.com/Joachricar/AliEC2/>

Project

AliEC2

- Provides virtual machines (VMs) for AliEN jobs.
 - CernVM used by default.
 - Job Agent started through a contextualisation script.
- VM interactions through a web-service interface.
 - Based on Dancer.
- Dedicated service for VM lifecycle management.
- Transparent to the end-user.

Project

- Web-service used to request/manage resources.
 - Dancer 2.
 - Easily configurable for use with reverse proxies.
 - Limits web-service access.
- VM-Pooling to reduce VM startup overhead.
 - Keeps a number of VMs alive after initial job execution.
 - Configurable max/min VM number.
 - Configurable idle timer before VM termination.

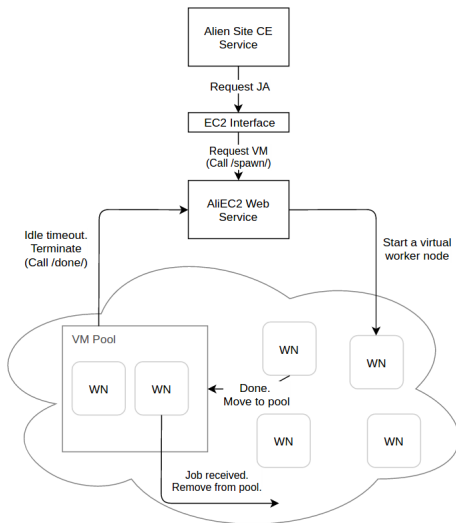


FIGURE – VM launch and lifecycle.

Current Status

- Fully working as intended.
 - VMs provided for job execution.
 - Minimal delay before job execution start.
 - Web-service access limited.
- Tested on both performance & robustness.
 - Initial results on performance are encouraging.
 - However...
 - Unknown delay upon submission of results.
 - Remains untested for 'real' AliEn jobs.
 - More testing needed.

Robustness

"The ability of a system to maintain an "acceptable" level of operability during exceptional, or unforeseen, circumstances".

Tests include :

- High load.
- Failures.
- Disruption of back-end services.

Initial results are promising.

Related Work

Likeness to an existing project – ElastiQ.²

- Manages VMs for job execution.
- Scales the VM number up/down based on present job number.
- Min/max VM quota.
- Transparent to end-users.

However :

- The presented project integrates directly with AliEn.
 - Scheduler/resource manager not required (e.g Torque/HTCondor).
 - VM Job Agents take jobs from the CS in a FIFO order.
- Web-service used for VM communications.
- Independent service for VM management.

2. <https://github.com/dberzano/elastiQ/>

Questions ?

Contact : sej001@student.uib.no

Source : <http://github.com/mtvedt/Sol/>