

An elastic OpenStack grid

Niklas Trippler

University of Bergen / Bergen University College
Supervisors – Bjarte Kileng, Kristin F. Hetland

What's in it for ALICE

- Ability to exploit external resources in periods with high demand
- Scale up and down applications on demand
 - Possibility of scaling up when resources are cheap
- Ability to quickly set up new grid sites
- Example: Grid site is running low on resources.
 - Deploy more worker nodes
- Example: AliEn's API servers are getting overloaded
 - Deploy more API servers

Goals for the project

- Automatic elasticity
- Scale up and down based on resource demand
 - Heat and Ceilometer
- Detecting situations when resources are running low and correcting
- Monitoring with ELK
 - ElasticSearch, Logstash and Kibana

Heat and Ceilometer

- Plugins for OpenStack
- Ceilometer collects data about resource usage
- Heat performs actions on that data
- Used here to measure resource usage and scale grid.


ELK

- Elasticsearch
 - High performance search and analytics tool
- Logstash
 - Input → filter → codec → output
- Kibana
 - Visualize the data

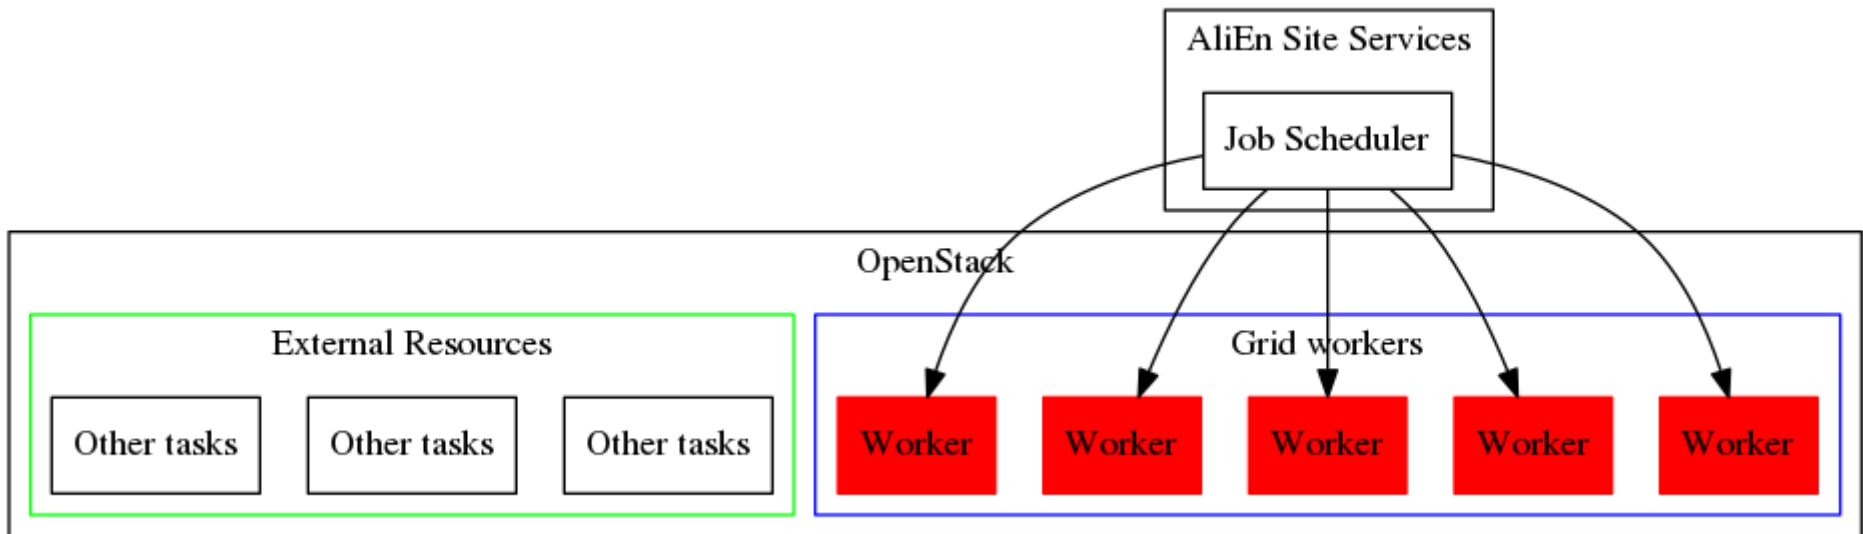
Secondary goals

- Recreatable environments via Puppet
 - Puppet is a tool for automatically recreating computer environments
 - Declare how the environment is supposed to look like, not how to do it
- ELK, OpenStack and integrations

Current solution [AliEn Grid site]



Thesis solution - Scale up on demand[Exploit external resources]



Current progress

- OpenStack environment up and running
- ELK almost up and running

Issues so far

- Puppet
- OS requirements
- Puppet modules
 - OpenStack
 - MySQL

In progress / Remaining work

- Deciding what grid tools to use for the prototype
- OpenStack monitoring metrics
- Getting grid and autoscaling implemented
- Measuring performance penalty of OpenStack

Outlook for end of project

- Able to set up a new grid site with minimal amount of work
- Grid able to scale with no, or minimal operator intervention
- Grid able to receive computational jobs
- The system able to give good metrics to operator via Kibana

Related Projects

- Elastiq from Dario Berzano
 - Monitors only batch queue
 - Only scaling layer
- AliEC2 improvement from Max at Bergen University College
 - Interface directly with AliEn
 - Replaces scheduler

Q&A

- Input and suggestions to niklas.trippler@cern.ch