

Elastic Grid computing on OpenStack

Sveinung Styve Rundhovde

University of Bergen / Western Norway University of Applied Science

Supervisors – Bjarte Kileng, Kristin F. Hetland

Project goal

- One-click deployment of AliEn Grid site on OpenStack
 - Recreatable environments with Heat Orchestration Templates and Puppet
- Automatic scaling based on resource usage
 - Integrate orchestration with telemetry to detect over/under utilization of resources and correcting
- Manual control
 - Adjust limits for automatic scaling or override
- Monitoring with ELK
 - Elasticsearch, Logstash and Kibana

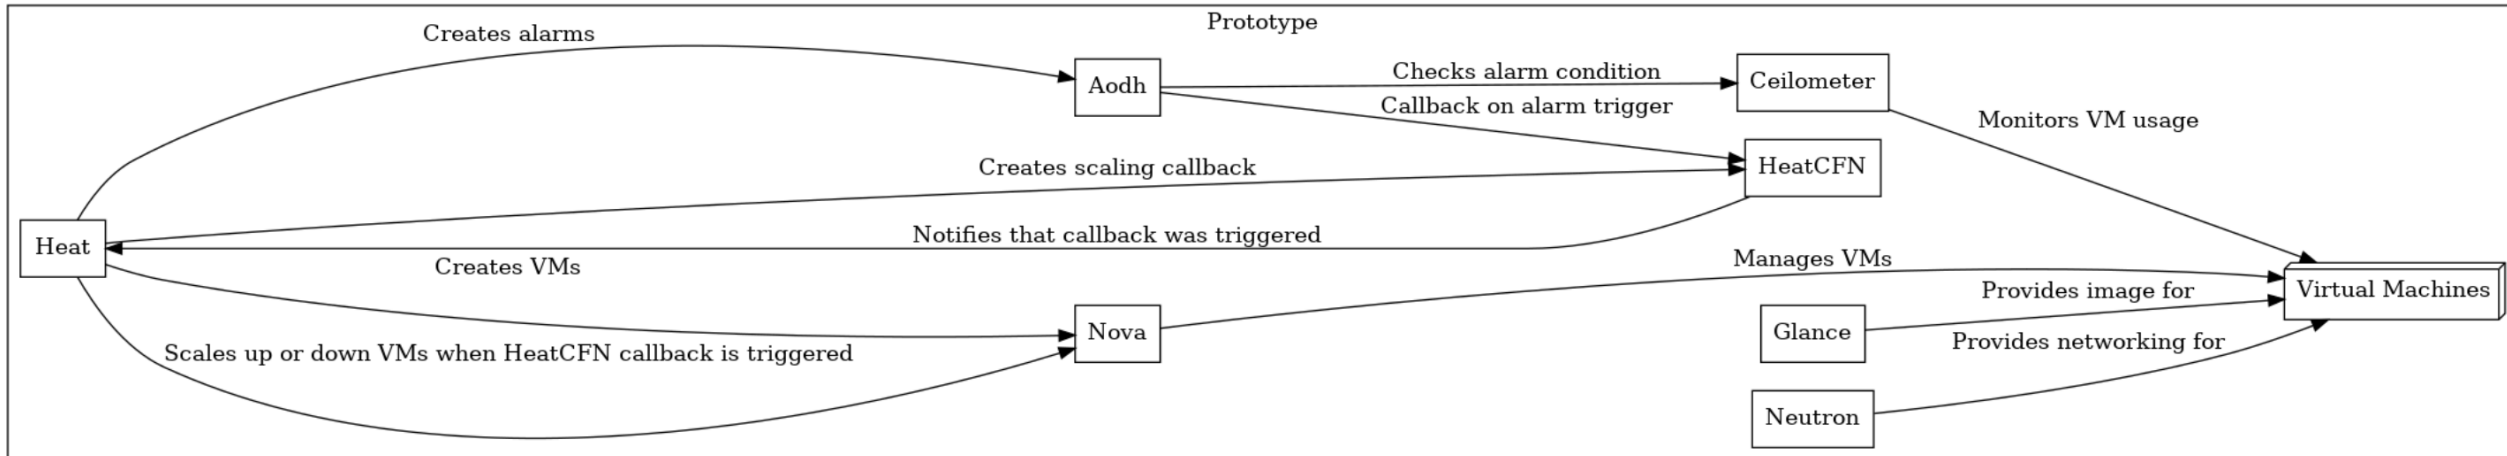
Benefits for ALICE

- Ability to quickly set up new Grid sites on OpenStack
 - Harness external resources in times with high demand
 - Possibility of scaling up when resources are cheap
- Less manual intervention
 - Resources scale automatically

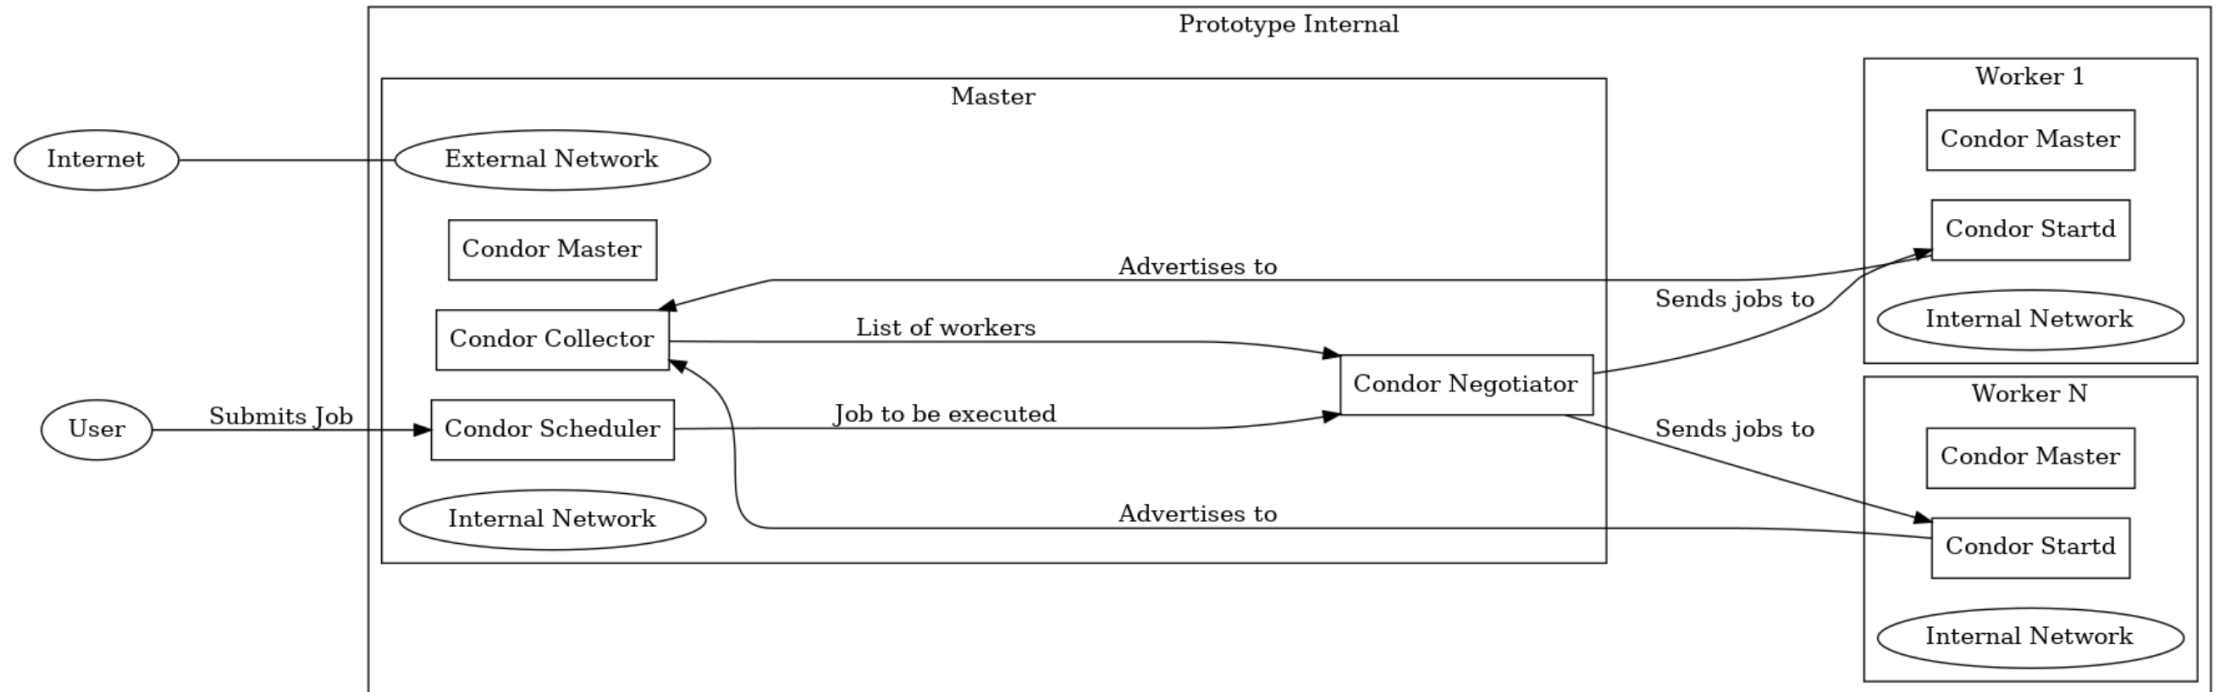
Current status

- Builds on an earlier project by Niklas Trippler
- Automatically deployed HTCondor batch system
- Scales automatically based on CPU usage
 - Upper and lower bounds can be set
 - Manual override is possible
- Resources and scaling policies defined in Heat Orchestration Templates (HOT)
 - Divided in inner and outer part for reusability
- Uses Heat and Ceilometer to set alarms and act when alarms are triggered (scale up/down)
- Worker nodes are removed in FIFO order, Heat does not let you decide which node to remove when scaling
- Some stability issues; creation and deletion of stack frequently fails, possibly due to unstable OpenStack environment set up with Packstack

Overview of external parts



Overview of internal parts



Suggested functionalities

- Senlin as an alternative to Heat?
 - Allows more fine grained control over stack
 - Possible to decide which VM to remove when scaling down
- Dashboard for setting parameters, launching, managing and monitoring stack
- Possibility of deploying a Squid server for CVMFS-caching
- Additional layer for handling communication with OpenStack
 - Should be able to discover and recover from OpenStack failures
- Set timers for scaling
 - Example use case: X VM instances has been acquired for Y hours; deploy X instances and take down after Y hours

Input and suggestions

- E-mail: sru083@student.uib.no