

## Why scheduling fairly?

Large public cloud



Small scientific datacentre



- approximately infinite resources
- tenants are charged *a posteriori*
- applications scale in/out freely
- saturated regime
- tenants are charged *a priori*
- advanced resource allocation needed

Default OpenNebula (ONE) scheduler is **FIFO** and based only on **static resources partitioning** among the projects → not suitable for scientific DC

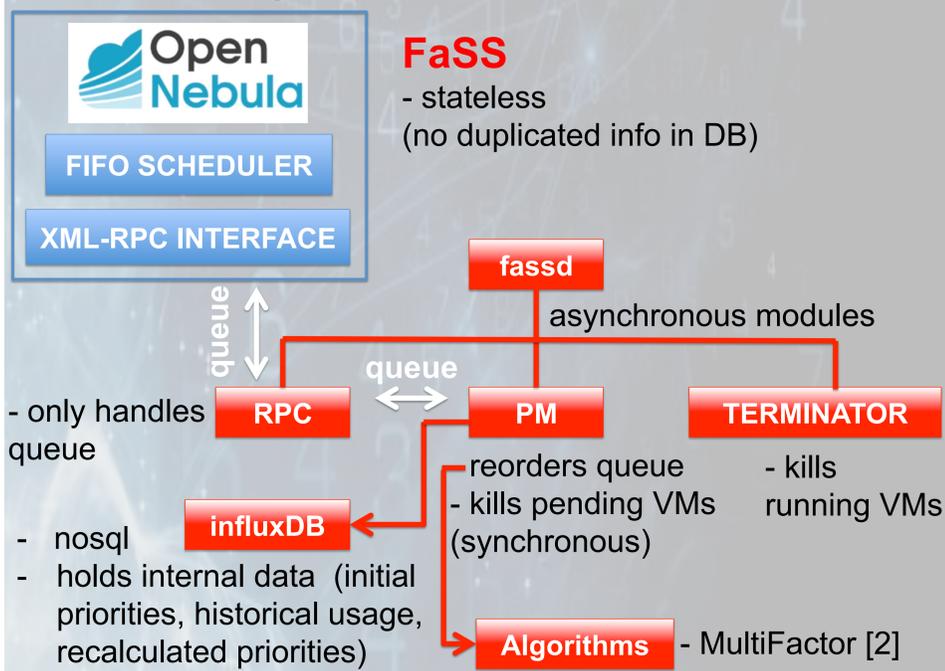
### FairShare Scheduling (FaSS) service [1]

Task priorities assigned according to

- initial weight
- historical resource usage

- **v1.0** released with ElectricIndigo in April 2017 with ONE patch
- **from v1.2** ONE 5.4 compatibility

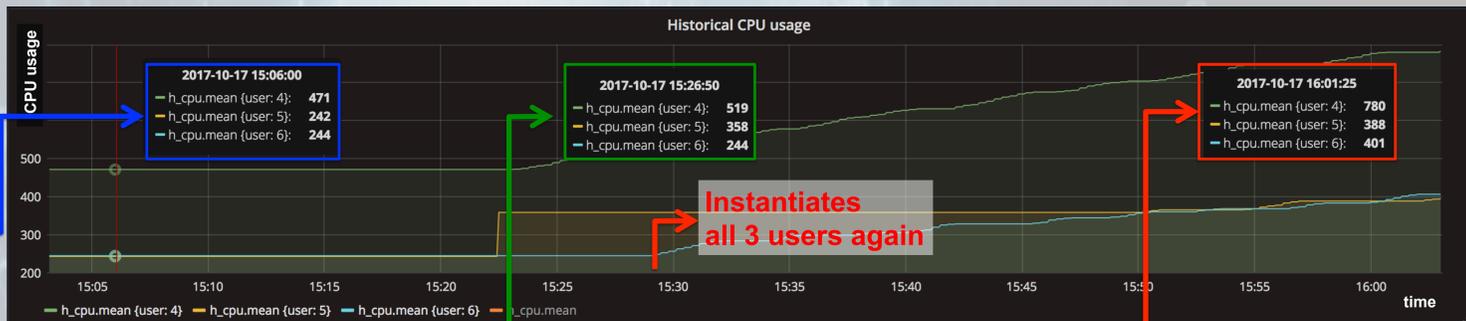
## High-level architecture



Test with three users, initial shares configured in `/etc/fass/shares.conf`

**Huey** uid=4 gid=1 share=50  
**Dewey** uid=5 gid=1 share=25  
**Louie** uid=6 gid=1 share=25

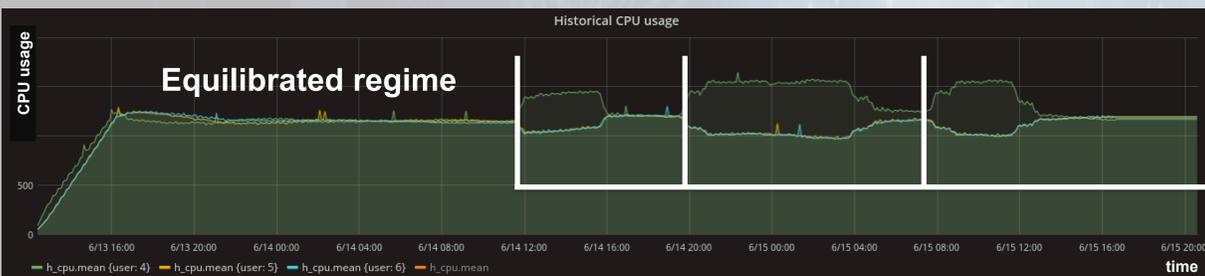
Periodically instantiating VMs for all 3 users



Stop instantiating Louie

Reaches configured shares

## Stress test



- Run for 1 week, ~22k VMs
- stable performance
- Adding periodical stress sources (0 – 300 VMs to random users)
- favours Huey
- after stress → goes back to equilibrium

## Installation

- Prerequisites:
  - install ONE (versions above 5.4 operate with FaSS above v1.2, previous versions of ONE need FaSS v1.1 or before)
  - install InfluxDB and create fassdb
- Install FaSS as root user:

```
yum localinstall one-fass-service-v1.4-1.4.x86_64.rpm
```
- Adjust the configuration file of the ONE scheduler, to allow it to point at the FaSS endpoint in `/etc/one/sched.conf`

## Usage

- Edit the initial shares for every user in `/tmp/one-fass/etc/shares.conf`
- Start FaSS: `systemctl start fass`

## Additional features

Keep your Cloud infrastructure clean with FaSS [3]

- Set VMs to be dynamic and to be terminated after a specific Time-to-Live (TTL)
- Instead of terminating, VMs can be powered-off, suspended or rebooted
- TTL settable per user.

## Outlook

- Testing new algorithms
- Integrating into the production infrastructure

Stay tuned!

References:

- [1] <https://github.com/indigo-dc/one-fass>
- [2] [https://slurm.schedmd.com/priority\\_multifactor.html](https://slurm.schedmd.com/priority_multifactor.html)
- [3] <https://opennebula.org/fass-fair-share-scheduler-for-opennebula/>

