

## Fair Share Scheduler for OpenNebula (FaSS): implementation and performance tests

*Tuesday, 10 July 2018 16:40 (20 minutes)*

A small Cloud infrastructure for scientific computing likely operates in a saturated regime, which imposes constraints to free applications' auto-scaling. Tenants typically pay a priori for a fraction of the overall resources. Within this business model, an advanced scheduling strategy is needed in order to optimize the data centre occupancy.

FaSS, a Fair Share Scheduler service for OpenNebula (ONE), addresses this issue by satisfying resource requests according to an algorithm, which prioritizes tasks according to an initial weight and to the historical resource usage of the project. In this contribution, we are going to describe the implementation of FaSS Version 1.0, released in March 2017 as a product of the INDIGO-DATA CLOUD project. The software was designed to be less intrusive as possible in the ONE code, and interacts with ONE exclusively through its XML-RPC interface. The native ONE scheduler is preserved for matching requests to available resources.

FaSS is made by five functional components: the Priority Manager (PM), a set of fair-share algorithms, Terminator (TM), the XML-RPC interface and the database. The main module, the PM, periodically requests the list of pending Virtual Machines (VMs) to ONE and re-calculates the priorities in the queue by interacting with an algorithm module of choice. In FaSS 1.0, the default algorithm is Slurm's MultiFactor. The TM module runs asynchronously with respect to the PM and it is responsible for removing from the queue VMs in pending state for too long as well as terminating or suspending running VMs after a configurable Time-to-Live. The XML-RPC server of FaSS intercepts the calls from the First-In-First-Out scheduler of ONE and sends back the reordered VMs queue. FaSS database is InfluxDB. It stores the initial and recalculated VM priorities and some additional information for accounting purposes. No information already present in the ONE DB is duplicated in FaSS.

In this contribution we are also going to show the results of FaSS functional and stress tests performed at the Cloud infrastructure of the INFN-Torino computing centre.

In this contribution we are also going to show the results of FaSS functional and stress tests performed at the Cloud infrastructure of the INFN-Torino computing centre.

**Primary authors:** Dr BAGNASCO, Stefano (Istituto Nazionale di Fisica Nucleare, Torino); VALLERO, Sara (Universita e INFN Torino (IT)); ZACCOLO, Valentina (Universita e INFN Torino (IT))

**Presenters:** Dr BAGNASCO, Stefano (Istituto Nazionale di Fisica Nucleare, Torino); VALLERO, Sara (Universita e INFN Torino (IT)); ZACCOLO, Valentina (Universita e INFN Torino (IT))

**Session Classification:** Posters

**Track Classification:** Track 7 –Clouds, virtualization and containers