

# Improved Cloud resource allocation: how INDIGO-Datacloud is overcoming the current limitations in Cloud schedulers

*Wednesday, 12 October 2016 11:15 (15 minutes)*

Performing efficient resource provisioning is a fundamental aspect for any resource provider. Local Resource Management Systems (LRMS) have been used in data centers for decades in order to obtain the best usage of the resources, providing their fair usage and partitioning for the users. In contrast, current cloud schedulers are normally based on the immediate allocation of resources on a first-come, first-served basis, meaning that a request will fail if there are no resources (e.g. OpenStack) or it will be trivially queued ordered by entry time (e.g. OpenNebula). This approach has been identified by the INDIGO-DataCloud project as being too simplistic for accommodating easily scientific workloads. Moreover, the scheduling strategies are based on a static partitioning of the resources, meaning that existing quotas cannot be exceeded, even if there are idle resources allocated to other projects. This is a consequence of the fact that cloud instances are not associated with a maximum execution time and leads to a situation where the resources are under-utilized. This is a non desirable situation in scientific data centers that struggle to obtain the maximum utilization of their resources. The INDIGO-DataCloud project is addressing the described issues in several different areas. On the one hand, by implementing fair-sharing strategies for OpenStack and OpenNebula through the “Synergy” and “FairShare Scheduler(FSS)” components, guaranteeing that the resources are accessed by the users according to the defined fair-share policies established by the system administrator. On the other hand, by implementing a mechanism to execute interruptible (or spot) instances. This way, higher priority instances (such as interactive nodes) can terminate lower priority instances that can be exploited by the users for fault-tolerant processing tasks. This way, it is possible to maximize the overall usage of an infrastructure (by filling the available resources with interruptible instances), without preventing users from running normal instances. Finally, taking into account that scientific data centers are composed of a number of different infrastructures (HPC, Grid, local batch systems, cloud resources), INDIGO is developing a “partition director” to address the problem of granting to a project the availability of a share quota of the physical resources in a center, while balancing their destination over different interface types, such as cloud and batch. This gives a resource provider the ability of dynamically resizing sub-quotas. This ability can also be transferred to the project, who can drive the resizing by controlling the resource request rate on the different infrastructures. This feature can be complemented by other optimizations implemented by INDIGO, such as the Synergy component. In this contribution, we will present the work done in the scheduling area during the first year of the INDIGO project in the outlined areas and the foreseen evolutions.

## Secondary Keyword (Optional)

Computing middleware

## Primary Keyword (Mandatory)

Cloud technologies

## Tertiary Keyword (Optional)

**Primary authors:** Dr LOPEZ GARCIA, Alvaro (Universidad de Cantabria (ES)); ZANGRANDO, Lisa (Università e INFN, Padova (IT)); SGARAVATTO, Massimo (Università e INFN, Padova (IT)); VALLERO, Sara (Università e INFN Torino (IT)); TANEJA, Sonia (Università e INFN, Bologna (IT)); BAGNASCO, Stefano (I.N.F.N. TORINO); DAL PRA, Stefano (INFN)

**Co-authors:** SALOMONI, Davide (Universita e INFN, Bologna (IT)); DONVITO, Giacinto (INFN-Bari)

**Presenter:** ZANGRANDO, Lisa (Universita e INFN, Padova (IT))

**Session Classification:** Track 7: Middleware, Monitoring and Accounting

**Track Classification:** Track 7: Middleware, Monitoring and Accounting