

Integrating cloud computing within an existing infrastructure

Monday, 16 September 2013 09:00 (8h 30m)

Description of Work

Starting from a virtualized infrastructure, the resources have been progressively integrated in a cloud environment powered by OpenStack.

Any new node coming to the datacenter is automatically provisioned (Razor) and configured (Puppet) as an Openstack Compute Node. This private cloud is partitioned in several pools (aggregate hosts), associated with different hardware types, each one dealing with different computing needs.

Therefore it is possible to dynamically shape and balance the resources allocated to the batch system (local and Grid), offer additional direct-cloud access to the users if required, spawn (and on-error migration) core services, etc.

Printable Summary

Cloud Computing is gaining interest within the scientific community. At IFCA datacenter, instead of providing a separated cloud computing service we have integrated it along with the rest of our infrastructure. Therefore we are able to offer the same set of resources, accessed through several interfaces according to the user needs: local batch system (Grid Engine), IaaS cloud (OCCI, OS, EC2) and Grid (UMD) access. At the same time, IT staff can use the cloud privately to spawn new services or balance the available resources easily between the different projects. This way, all the services and computing nodes are spawned on top of this cloud infrastructure.

The flexibility obtained from this cloud-based infrastructure makes possible to satisfy computing demands not fitting in the previous old-static schema.

Primary authors: Mr LOPEZ GARCIA, Alvaro (CSIC); ORVIZ, Pablo (CSIC)

Co-authors: FERNANDEZ, Enol (CSIC); Dr PLASENCIA, Isabel (CSIC)

Presenters: Mr LOPEZ GARCIA, Alvaro (CSIC); ORVIZ, Pablo (CSIC)

Session Classification: Posters display