

The 'Cloud Area Padovana': lessons learned after two years of a production OpenStack-based IaaS for the local INFN user community

Friday, March 10, 2017 9:00 AM (30 minutes)

The Cloud Area Padovana is an OpenStack-based scientific cloud, spread across two different sites - the INFN Padova Unit and the INFN Legnaro National Labs - located 10 km away but connected with a dedicated 10 Gbps optical link. In the last two years its hardware resources have been scaled horizontally by adding new ones: currently it provides about 1100 logical cores and 50 TB of storage.

Special in-house developments were also integrated in the OpenStack dashboard, such as a tool for user and project registrations with direct support for Single Sign-On via the INFN-AAI Identity Provider as a new option for the user authentication. The collaboration with the EU-funded INDIGO-DataCloud project, started one year ago, allowed to experiment the integration of Docker-based containers and the fair-share scheduling: a new resource allocation mechanism analogous to the ones available in the batch system schedulers for maximizing the usage of shared resources among concurrent users and projects. Both solutions are expected to be available in production soon.

The entire computing facility now satisfies the computational and storage demands of more than 100 users afferent to about 25 research projects. In this paper we'll present the architecture of the Cloud infrastructure, the tools and procedures used to operate it ensuring reliability and fault-tolerance. We'll especially focus on the lessons learned in these two years, describing the challenges identified and the subsequent corrective actions applied. From the perspective of scientific applications, we'll show some concrete use cases on how this Cloud infrastructure is being used. In particular we'll focus on two big physics experiments which are intensively exploiting this computing facility: CMS and SPES. CMS deployed on the cloud a complex computational infrastructure, composed of several user interfaces for job submission in the Grid environment/local batch queues or for interactive processes; this is fully integrated with the local Tier-2 facility. To avoid a static allocation of the resources, an elastic cluster, based on cernVM, has been configured: it allows to automatically create and delete virtual machines according to the user needs. SPES is using a client-server system called TraceWin to exploit INFN's virtual resources performing a very large number of simulations on about a thousand nodes elastically managed.

Summary

In this paper we'll present the architecture of the Cloud infrastructure, the tools and procedures used to operate it ensuring reliability and fault-tolerance. We'll especially focus on the lessons learned in these two years, describing the challenges identified and the subsequent corrective actions applied. From the perspective of scientific applications, we'll show some concrete use cases on how this Cloud infrastructure is being used. In particular we'll focus on two big physics experiments which are intensively exploiting this computing facility: CMS and SPES.

Primary authors: Mr CRESCENTE, Alberto (INFN-Padova); Dr DORIGO, Alvise (INFN-Padova); Dr KONOMI, Ervin (INFN-Padova); Dr CHIARELLO, Fabrizio (INFN-Padova); Dr FANZAGO, Federica (INFN-Padova); Mrs COSTA, Fulvia (INFN-Padova); Dr ZANGRANDO, Lisa (INFN-Padova); Dr VERLATO, Marco (INFN-Padova); Dr SGARAVATTO, Massimo (INFN-Padova); Dr SEGATTA, Matteo (INFN-Padova); Dr TRITTO, Nicola (INFN-Padova); Dr ANDREETTO, Paolo (INFN-Padova); Dr FANTINEL, Sergio (INFN-LNL); Dr TRALDI, Sergio (INFN-Padova)

Presenter: Dr VERLATO, Marco (INFN-Padova)

Session Classification: Infrastructure Clouds and Virtualisation I

Track Classification: Infrastructure Clouds and Virtualisation