

# The Cloud Area Padovana: from pilot to production

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

The Cloud Area Padovana has been running for almost two years. This is an OpenStack-based scientific cloud, spread across two different sites: the INFN Padova Unit and the INFN Legnaro National Labs.

The hardware resources have been scaled horizontally and vertically, by upgrading some hypervisors and by adding new ones: currently it provides about 1100 cores.

Some in-house developments were also integrated in the OpenStack dashboard, such as a tool for user and project registrations with direct support for the INFN-AAI Identity Provider as a new option for the user authentication.

In collaboration with the EU-funded Indigo DataCloud project, the integration with Docker-based containers have been experimented and will be available in production soon.

This computing facility now satisfies the computational and storage demands of more than 70 users afferent to about 20 research projects.

We present here the architecture of this Cloud infrastructure, the tools and procedures used to operate it. We also focus on the lessons learnt in these two years, describing the problems that were found and the corrective actions that had to be applied. We also discuss about the chosen strategy for upgrades, which combines the need to promptly integrate the OpenStack new developments, the demand to reduce the downtimes of the infrastructure, and the need to limit the effort requested for such updates.

We also discuss how this Cloud infrastructure is being used. In particular we 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, using a client-server system called TraceWin, exploits INFN's virtual resources performing a very large number of simulations on about a thousand nodes elastically managed.

## Secondary Keyword (Optional)

Computing facilities

## Primary Keyword (Mandatory)

Cloud technologies

## Tertiary Keyword (Optional)

**Primary authors:** CRESCENTE, Alberto (Universita e INFN, Padova (IT)); DORIGO, Alvise (Universita e INFN, Padova (IT)); FANZAGO, Federica (Universita e INFN, Padova (IT)); COSTA, Fulvia (INFN Padova); ZANGRANDO, Lisa (Universita e INFN, Padova (IT)); VERLATO, Marco (Universita e INFN, Padova (IT)); SGARAVATTO, Massimo (Universita e INFN, Padova (IT)); ANDREETTO, Paolo (INFN Padova); FANTINEL, Sergio (INFN Laboratori Nazionali di Legnaro); TRALDI, Sergio (INFN Padova)

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

**Session Classification:** Track 3: Distributed Computing

**Track Classification:** Track 3: Distributed Computing