

# Exploiting private and commercial clouds to generate on-demand CMS computing facilities with DODAS

Monday 9 July 2018 12:00 (15 minutes)

Reducing time and cost, through setup and operational efficiency increase is a key nowadays while exploiting private or commercial clouds. In turn this means that reducing the learning curve as well as the operational cost of managing community-specific services running on distributed environments became a key to success and sustainability, even more for communities seeking to exploit opportunistic cloud resources. DODAS (Dynamic On Demand Analysis Service) is a tool based on INDIGO-DataCloud services which allows on-demand generation of a container-based HTCondor batch system over cloud infrastructures implementing a “Batch System as a Service” that can be easily customized to cover multiple scenarios. DODAS has been designed to completely automate the process of provisioning, creating, managing and accessing a pool of heterogeneous computing and storage resources over cloud infrastructures. As such, DODAS provides a complete abstraction of the underlying IaaS, which is only required to provide VMs or bare metal resources. It uses a combination of TOSCA templates and Ansible roles to orchestrate and configure cloud resources, and it relies on Apache Mesos to abstract CPU, RAM and storage. Marathon is adopted as container orchestration platform on top of Mesos. Marathon takes care of setting up both HTCondor and any additional services (e.g squid proxy, proxy certificate cache) that experiments might require. Credential management is completely Json Web Tokens-based which relies on Identity Access Management (IAM) and Token Translation Services (TTS) provided by INDIGO. The DODAS architecture is highly flexible and customizable; as such, it has been evolved in order to provide also the configuration and the automatic setup services and dependencies required by a regular CMS computing facility, a Tier site. This means that DODAS allows CMS to utilize “any cloud provider” to generate sites on demand, with almost zero effort. We describe the overall architecture of DODAS, including the integration strategy adopted to connect it with the CMS HTCondor Global Pool. We also present performance, scalability and efficiency results of DODAS-generated Tiers used to process CMS jobs, showing how DODAS has been effectively used in different CMS sites in order to consolidate disparate computing resources being used by CMS scientists. We will present the different ways DODAS can be used to integrate different usage patterns, especially how it can fit the final stages of the CMS analysis.

Finally, we present and compare results obtained both on private and commercial clouds spanning from Microsoft Azure, to Open Telekom Cloud up to private OpenStack instances.

**Author:** SPIGA, Daniele (Universita e INFN, Perugia (IT))

**Co-authors:** Mr LOPEZ GARCIA, Alvaro (CSIC); PALACIDO, aida (ifca); DUMA, Cristina (INFN-CNAF); Mr TRACOLLI, Mirco (INFN-Perugia); BOCCALI, Tommaso (INFN Sezione di Pisa, Universita' e Scuola Normale Superiore, P); CIANGOTTINI, Diego (Universita e INFN, Perugia (IT)); ANTONACCI, Marica; CECCANTI, Andrea; SALOMONI, Davide (Universita e INFN, Bologna (IT)); DONVITO, Giacinto (INFN-Bari)

**Presenter:** SPIGA, Daniele (Universita e INFN, Perugia (IT))

**Session Classification:** T7 - Clouds, virtualization and containers

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