

DEEP-HybridDataCloud

Project overview, achievements

A. Costantini, C. Duma

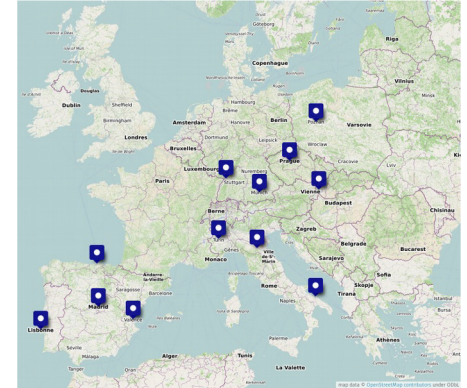
2/10/2020

Outline

- **DEEP-Hybrid DataCloud overview**
- **Overview of technical activities**
- **Conclusions**

DEEP project in a nutshell

- **Designing and Enabling E-Infrastructures** for intensive data Processing in a **Hybrid DataCloud** (Grant agreement number 777435, Nov 2017 – Apr 2020)
 - 10 Partners, 5 Countries
- **Global objective:** Promote the use of **intensive computing services** by different research communities and areas, and the **support by the corresponding e-Infrastructure** providers and open source projects
 - Focusing on **Machine learning, Deep learning, and Post processing**



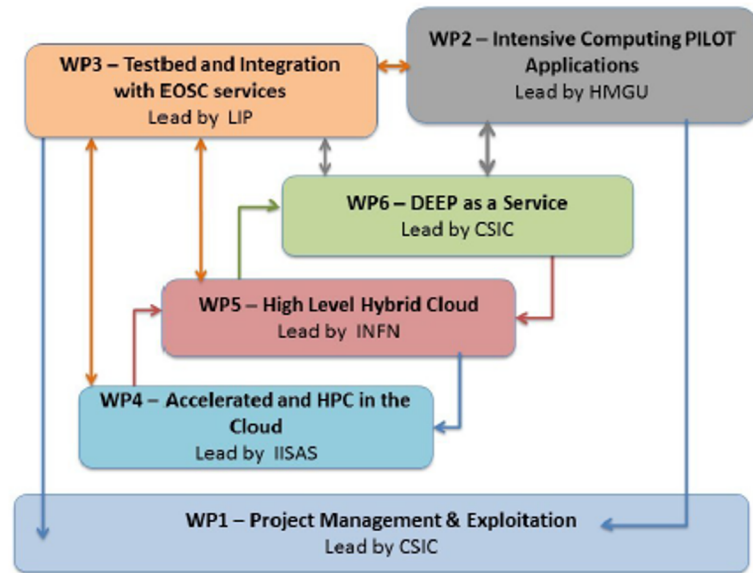
DEEP-HybridDataCloud vision

- Provide **specialized cloud services** to **develop, exploit and share** machine learning and deep learning applications → **Service Oriented Architectures and platforms**
 - Covering the whole application development cycle
 - Focused on all types of users
- **Transparent access** to specialized computing resources (accelerators, high performance computing) → **reduce entry barrier**
- Build the **EOSC machine learning marketplace** as an application exchange → **ease of use, foster collaboration, knowledge dissemination**
- Using different resources (Cloud, HPC, GPU, etc.)
→ **ease of use**

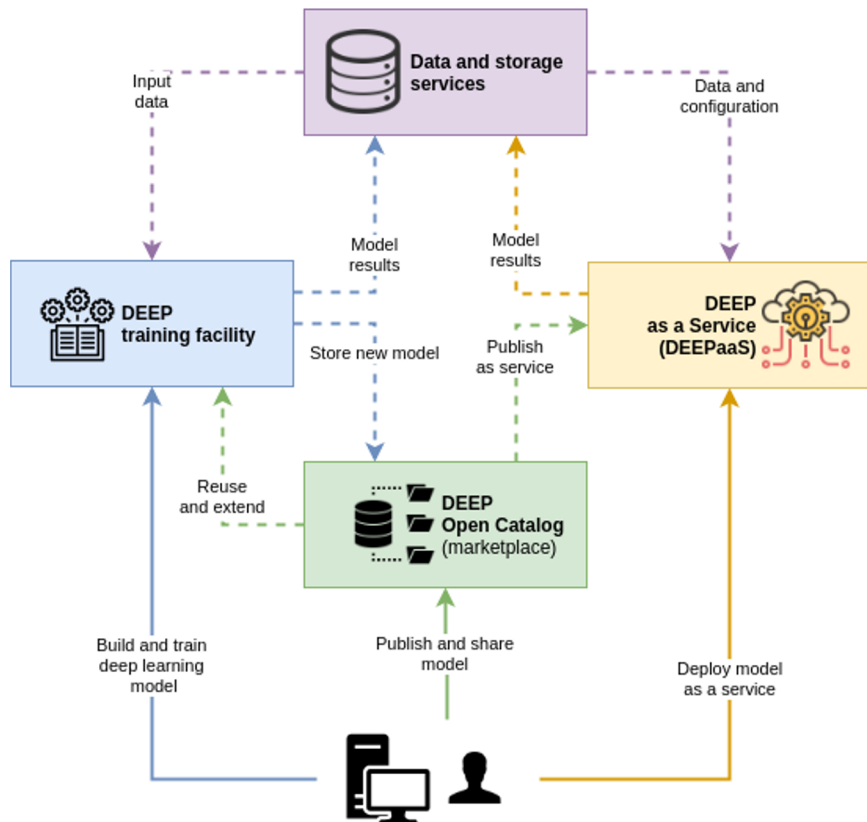


DEEP WP Breakdown

- **WP1 (NA):** Project Management and Exploitation
 - Project oversight, quality management, admin, etc.
- **WP2 (NA):** Intensive Computing Pilot Applications
 - Definition and understanding of pilot usage scenarios
- **WP3 (SA):** Testbed and Integration with EOSC service
 - Project service activities: testbeds, tools, integrations
- **WP4 (JRA):** Accelerated High Performance Computing in the Cloud
 - Support for accelerators and HPC resources
- **WP5 (JRA):** High Level Hybrid Cloud solutions
 - Platform provisioning, delivering the execution platform
- **WP6 (JRA):** DEEP as a Service
 - Deliver final solution to the users
- **WP7:** Ethics requirements
 - Leader: CSIC

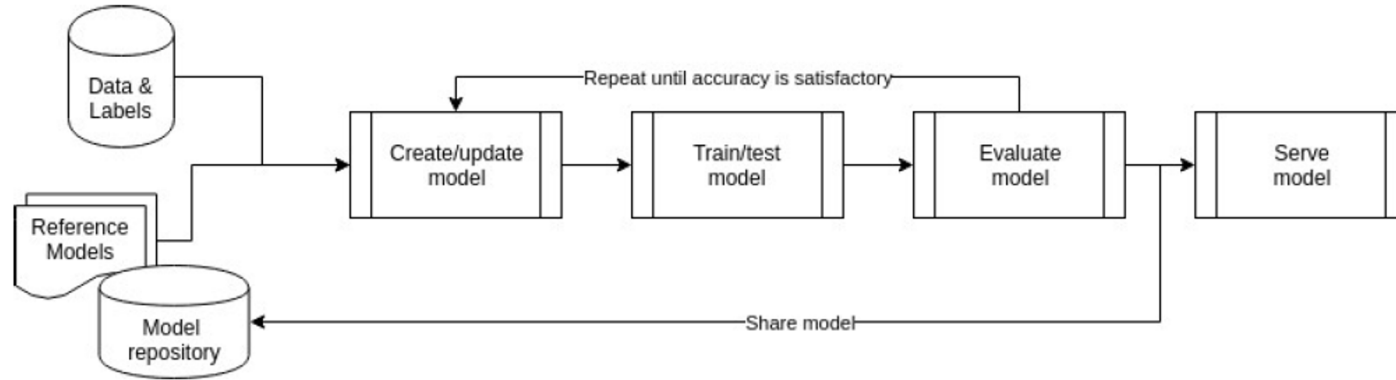


High level architecture



- Position as technology providers to support DL/ML in the EOSC
- Generic building blocks (services) for exploitation through EOSC
 - **DEEP training facility**
 - **DEEP as a Service (DEEPaaS)**
 - **DEEP Open Catalog**
- Integration with storage from external initiatives (eXtreme-DataCloud, EGI Data Hub)

Machine learning development cycle


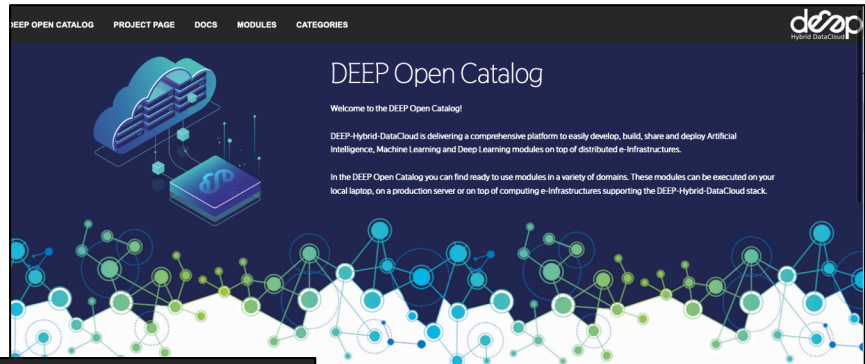
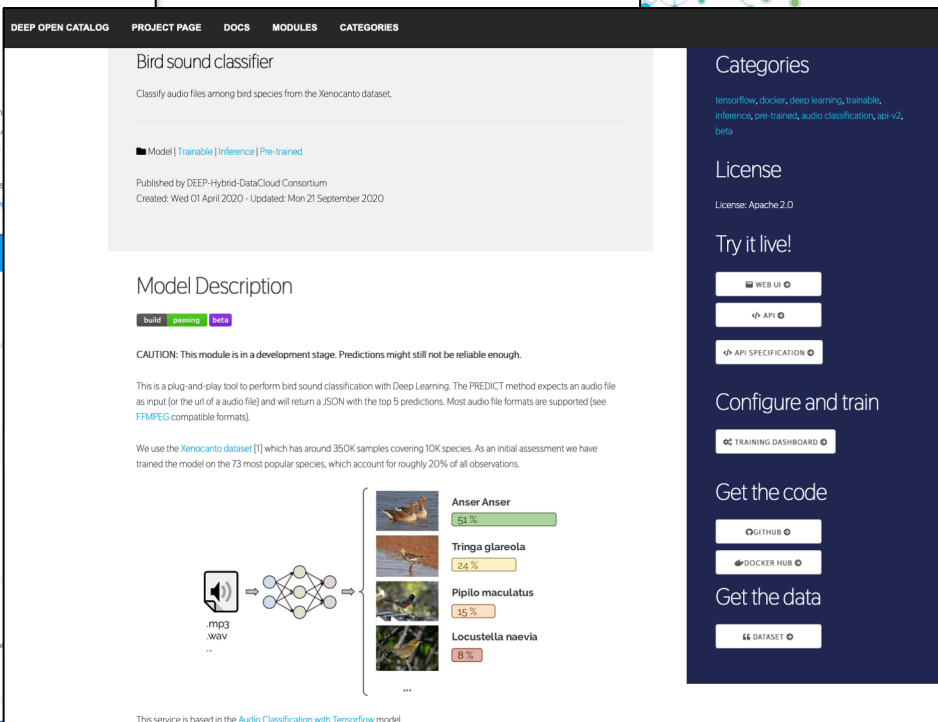
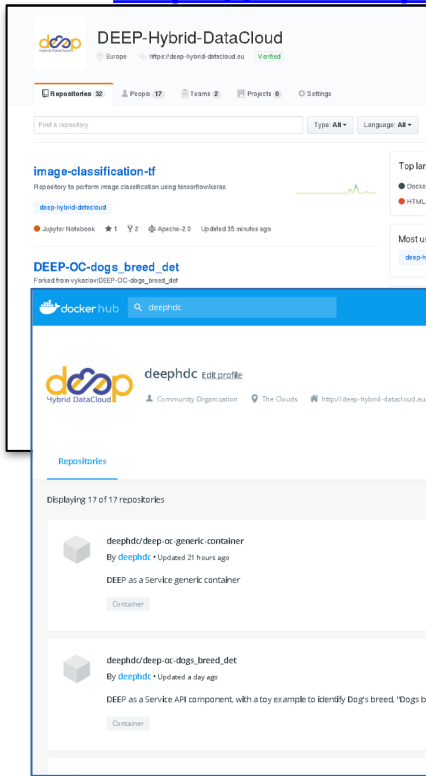


- The project covers all development phases of a machine learning application:
 - Creation, update, improvement of a model (existing or new) → **DEEP Open Catalog and DEEP training facility**
 - Model training, test, evaluation → **DEEP training facility**
 - Model deployment as a service → **DEEP as a Service**
 - Model publication for sharing and reuse → **DEEP Open Catalog**

DEEP Open Catalog

Sharing knowledge between users

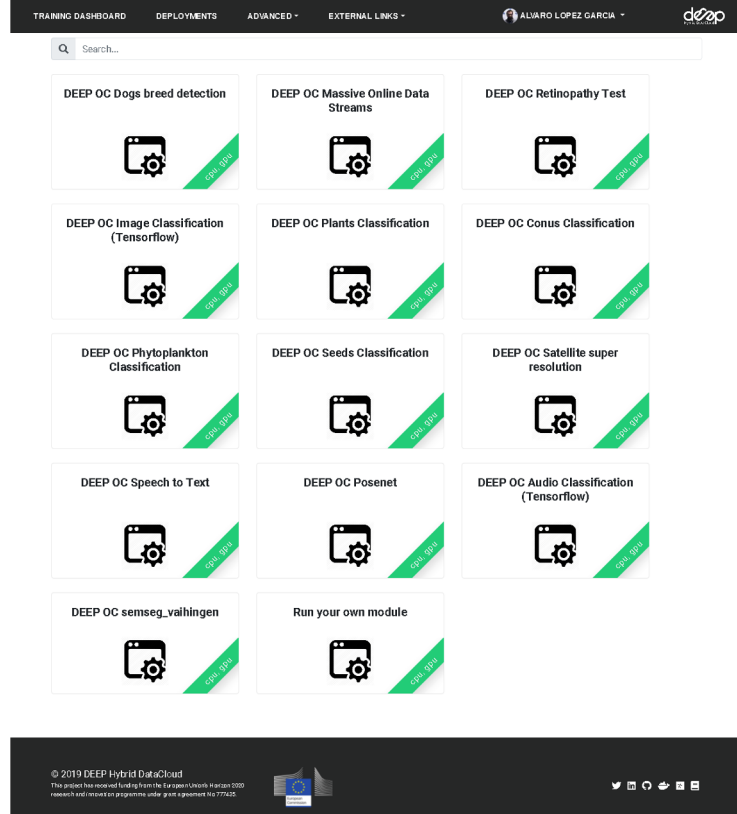
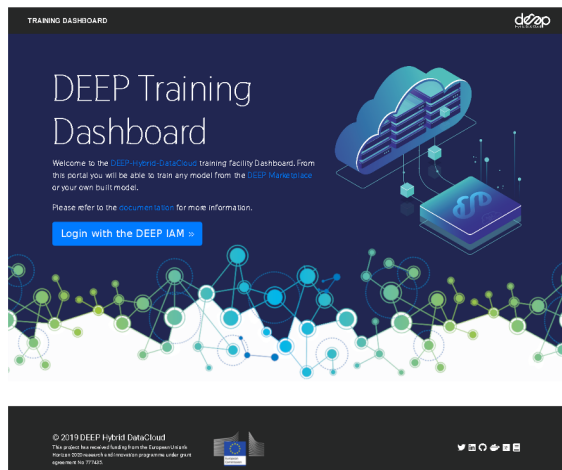
<https://marketplace.deep-hybrid-datacloud.eu>



DEEP training facility

Development and training environments

<https://train.deep-hybrid-datacloud.eu>



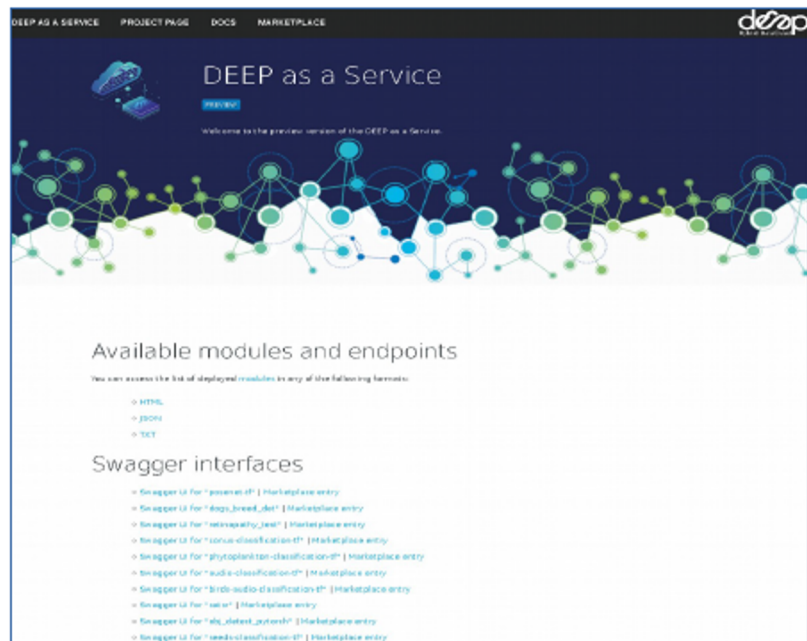
Train, test and evaluate of models through the DEEPaaS API

- Transparent access to infrastructure resources
- Based on **TOSCA templates** submitted via **INDIGO-DataCloud PaaS orchestrator** and **Infrastructure Manager**

DEEPaaS: DEEP as a Service

Automatic deployment of catalog modules to exploit their functionality

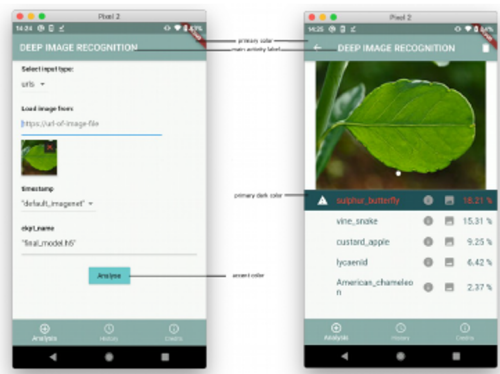
<https://deepaas.deep-hybrid-datacloud.eu>



DEEPaaS API: models as services

Exposing models through a common and standard-based API

<https://deepaas.readthedocs.io>



Deep Learning Applications

INTRO IMAGE RECOGNITION MISCELLANEOUS

This webpage gathers all the applications developed at the Instituto de Física de Cantabria (IFCA) using deep learning techniques.

Image Recognition



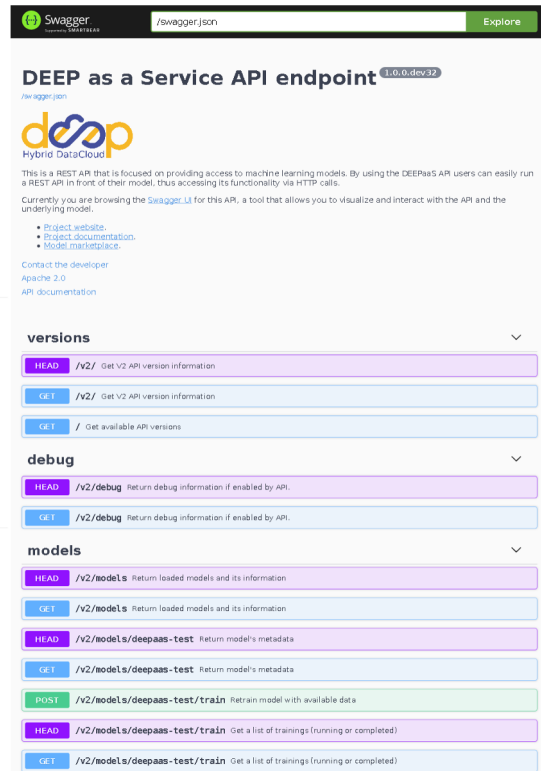
Plants

TRY IT!

Author Ignacio Heredia

Description

This application classifies a plant image, returning the most likely species from a list of Western European.



Costantini A., Duma C. - DEEP-HDC project overview

Different users, different needs

- **Category 1:** Deploy a readily trained network for somebody else to use it on his/her data set
 - Domain knowledge
- **Category 2:** Retrain (parts of) a trained network to make use of its inherent knowledge and to solve a new learning task
 - Domain + machine learning knowledge
- **Category 3:** Completely work through the deep learning cycle with data selection, model architecture, training and testing
 - Domain + machine + technological knowledge

High Level Hybrid PaaS

- Advance TRL of INDIGO-DataCloud PaaS Orchestrator, Infrastructure Manager, Virtual Router and Cloud Info Provider in order to:
 - Scale deployments across private and public clouds with seamless network connections
 - Exploit special hardware resources (namely GPU, Infiniband) from Cloud and HPC environments
 - Ensure flexibility and interoperability through TOSCA and federated authentication (IAM)
- Improved support for hardware resource selection
 - Fine-grained accelerator (GPU) selection
 - Support at information providers modules
- Improved support for hybrid-cloud deployments
 - Leveraging high-level networking solutions
- Improved management of deployment failures
 - Resource selection and error recovery

Overview of technical activities

CNAF for the DEEP-HDC project

- 3M euros (Total project Funding)
 - 375K (BARI, CNAF, TORINO)
 - 110K Personale (CNAF)
 - 20K Missioni (Progetto, CNAF_CCR)

- **448 PMs** over 30 months
 - 56 PMs (INFN)
 - 22 PMs (CNAF)

WP	Task	PM complessivi	CNAF (PM)	Reponsabilità
WP3		12		
WP3	Task 3.1 – Pilot testbeds and integration with EOSC platform and their services	4	3	
WP3	Task 3.2 – Software quality assurance, release, maintenance and support	8	8	Responsabilità INFN: Cristina Duma
WP4		12		
WP4	Task 4.3 – Interaction with HPC resources with PaaS approach	8	1	
WP5	INFN WP Leader	20		WP Leader: Giacinto Donvito/Marica Antonacci
WP5	Task 5.1 – PaaS-level Orchestration Supporting Multi-IaaS Hybrid Infrastructures	12	3	Responsabilità INFN: Marica Antonacci
WP5	Task 5.3 – High-level networking orchestration to connect seamlessly to hybrid clouds	3	3	
WP6		12		
WP6	Task 6.1 – Application model building	8	3	
WP6	Task 6.2 – DEEP as a Service deployment and exploitation	4	1	
Totale		56	22	

CNAF: people who reported

Persone	%	WP	Mesi	PM
Salomoni D.	10%	3,5,6	28	2
Duma C.	20%	3,5,6	28	5
Michelotto D.	7%	3	27	2
Ceccanti A.	7%	4,5	28	2
Ronchieri E.	7%	3,6	28	2
Costantini A.	20%	3,4,5,6	21	5
Morganti L.	15%	2,4,5,6	17	3
Falabella A.	7%	5,6	27	2
Vianello E.	20%	3,4	14	2
Caberletti M.	0%	5	0	0

WP3 - Pilot infrastructures and services

LIP/INCD (Provider)

IM: <https://im.ncg.ingrid.pt:8800>
Alien4Cloud: <https://a4c.ncg.ingrid.pt>
OpenStack: <https://stratus.ncg.ingrid.pt/>

CESNET (Provider)

OpenStack: <https://dashboard.cloud.muni.cz/>

IFCA/CSIC (Provider)

DEEPaaS: <https://deepaas.cloud.ifca.es/api/v1/web/deepaas>
Training dashboard: <https://train.deep-hybrid-datacloud.eu/>
Mesos (GPUs): <https://mesos.cloud.ifca.es/mesos/>
OpenStack (DEEP-IAM): <https://portal.cloud.ifca.es/>
Oneprovider (DEEP-IAM): <https://oprov.ifca.es>
Object store: Swift: <https://cephrgw01.ifca.es:8080/swift/v1>

PSNC (Provider)

QCG/HPC: `ssh ui.eagle.man.poznan.pl`
Mesos: <https://cereus.man.poznan.pl/mesos>
OpenStack: <https://cereus.man.poznan.pl/horizon>

IISAS (Provider)

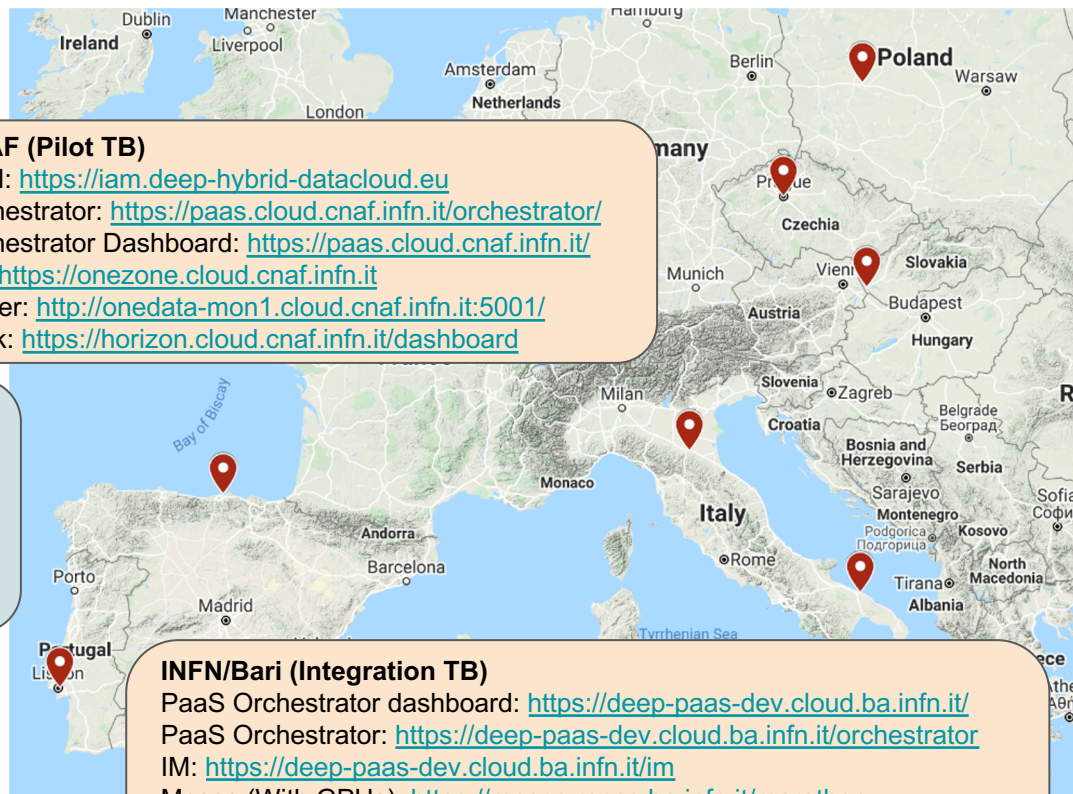
Mesos (GPUs): <https://mesos.ui.sav.sk/mesos-web/>
OpenStack (DEEP-IAM): <https://horizon.ui.savba.sk/horizon/>
Oneprovider (DEEP-IAM): <https://oneprovider.fedcloud.eu>

INFN/CNAF (Pilot TB)

DEEP-IAM: <https://iam.deep-hybrid-datacloud.eu>
PaaS Orchestrator: <https://paas.cloud.cnaf.infn.it/orchestrator/>
PaaS Orchestrator Dashboard: <https://paas.cloud.cnaf.infn.it/>
Onezone: <https://onezone.cloud.cnaf.infn.it>
Oneprovider: <http://onedata-mon1.cloud.cnaf.infn.it:5001/>
OpenStack: <https://horizon.cloud.cnaf.infn.it/dashboard>

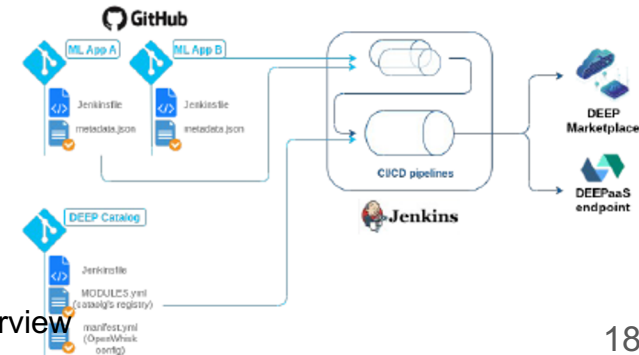
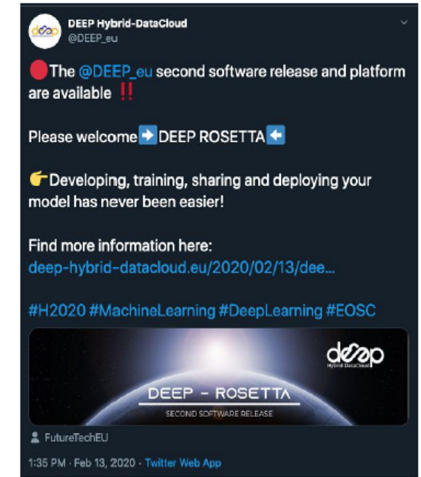
INFN/Bari (Integration TB)

PaaS Orchestrator dashboard: <https://deep-paas-dev.cloud.ba.infn.it/>
PaaS Orchestrator: <https://deep-paas-dev.cloud.ba.infn.it/orchestrator>
IM: <https://deep-paas-dev.cloud.ba.infn.it/im>
Mesos (With GPUs): <https://mesos.recas.ba.infn.it/marathon>
Onezone: <https://onezone-beta.cloud.ba.infn.it>
Oneprovider: <https://oneprovider-beta.cloud.ba.infn.it>
Vault: <https://vault.cloud.ba.infn.it:8200>
OpenStack: <https://cloud.recas.ba.infn.it>



WP3 - Software quality assurance, release, maintenance and support

- Release management
 - Software and platform release and maintenance
 - Project-wide documentation, product documentation, release notes
- Security coordination and external assessment
- Extensive automation in place for relevant processes
 - Software Quality Assurance (SQA), Software release and maintenance
 - State of the art continuous integration infrastructure
- Work with user communities with DevOps approach
 - Collaboration with WP2



WP4 - Interaction with HPC resources with PaaS approach

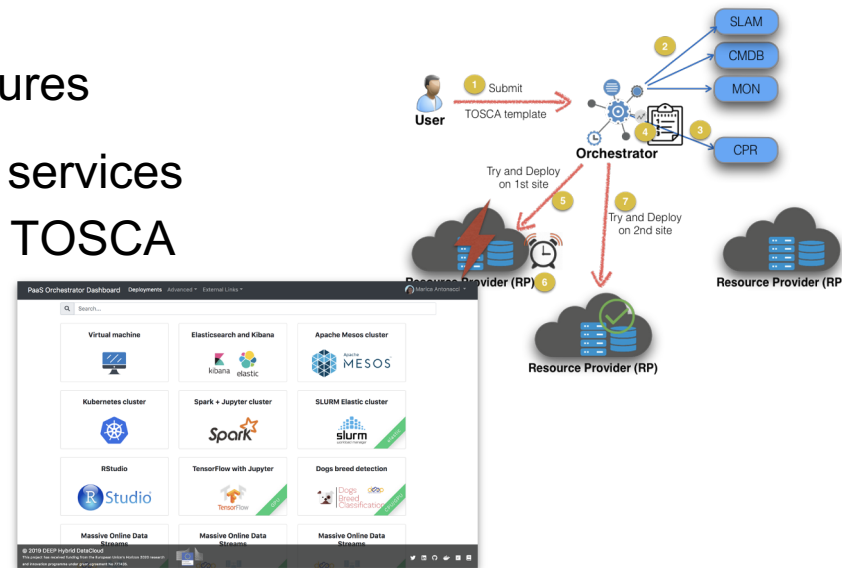
Deployment of a Kubernetes cluster

- Support for GPU in Kubernetes is provided via NVIDIA runtime plugin
- Integration with DEEP IAM using OpenID Connect (OIDC) standard
- Authorization supported via Attribute Based Access Control (ABAC) in Kubernetes
- Support to the deployment of production site with GPUs and OIDC

WP5 - PaaS-level Hybrid Cloud solutions

5.1 Support to Multi-IaaS Hybrid Infrastructures

- Integration tests with the developed services
- Testing of automatic deployment via TOSCA
- Orchestrator-Dashboard
 - tests and deployments



5.3 High-level networking orchestration

- This activity has not been finalised due to people leave

WP6 - DEEP as a Service

- Alien4Cloud high level user interface as DEEPaaS solution
- Application lifecycle contribution
 - deployment, testing and validation

Alien4Cloud
(Application
Modeler)

tosca_definitions_version: tosca_simple_yaml_1_0

imports:

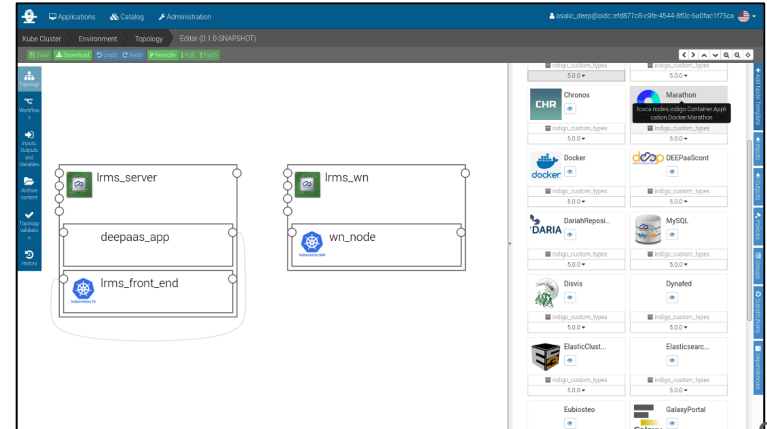
- indigo_custom_types: https://raw.githubusercontent.com/indigo-dc/tosca-types/master/custom_types.yaml

topology_template:

inputs:

wn_num:
type: integer
description: Number of WNs in the cluster
default: 1
required: yes

...
wn_node:
type: tosca.nodes.indigo.LRMS.WorkerNode.Kubernetes
properties:
front_end_ip: { get_attribute: [lrms_server, private_address, 0] }
requirements:
- host: lrms_wn



Conclusions

- DEPP-Hybrid DataCloud outcomes and results presented at the final review
 - Held on June 4th 2020, virtually
 - Excellent result in the final evaluation
 - *“Project has delivered exceptional results with significant immediate or potential impact”*
 - CNAF had an important role in WP3 (Release management and Software Quality Control)
 - This central role and the related activities have been recognized by the reviewers
 - *“WP3 [...] it ensured the production of an industrial quality software”*
- CNAF contribution
 - Depending on the effort distribution, CNAF supported all the JRAs by contributing to deployment, testing and validation of the tools and services
 - Used in the current e-infrastructures
 - Available in the EOSC marketplace
- Other contribution
 - Fund fixed-term contracts
 - Strengthen the relation with other players and partners