

PaaS Solutions TCOM Area

Marica Antonacci INFN



Dissemination level: Public/Confidential *If confidential, please define:* Disclosing Party: (those disclosing confidential information) Recipient Party: (to whom this information is disclosed, default: project consortium)



EOSC-hub receives funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 777536.



PaaS solutions TCOM Area

- Macro-features and functionalities
- Reference Architecture
- Standards & Protocols
- Interoperability
- Open points and next steps

EOSC-hub PaaS Solutions macro-features (1)

Advanced Cloud Services Orchestration

- Abstract the laaS resources
 - Automatic discovery of cloud compute/storage providers, capabilities and service availability
 - Seamless access to diverse computing environments
 - Private clouds (OpenStack, OpenNebula, etc.)
 - Public clouds (Amazon, Azure, Google)
 - HPC sites
- Manage the provisioning, configuration and scaling of the cloud resources
- Support Hybrid deployments

[Offered by INDIGO Orchestrator + IM + CLUES + vRouter]

- Support for network orchestration
 - Connect resources on different sites creating overlay networks

EOSC-hub PaaS Solutions macro-features (2)

- Advanced Scheduling capabilities
 - *Data-aware scheduling*: take into account the data location when dispatching the deployment requests
 - Special hw requirements (GPU, FPGA, Infiniband)
 - *Retry strategy*: provide a mechanism to address the deployment failures rescheduling the deployment to another site
 [Offered by INDIGO-DC PaaS Orchestrator]
- Support for data orchestration and movements
 - Provide policy-driven data management
 - Handling QoS for data replicas

[Will be offered by INDIGO Orchestrator + Rucio]

EOSC-hub PaaS Solutions macro-features (3)

• Provide solution stack as a Service

[Offered by INDIGO Orchestrator + IM]

- Support for container-based infrastructures and cloud native stacks
 - Complex topologies such as Kubernetes cluster, Mesos cluster, HTCondor cluster
- Provide on-demand application deployment
 - Start long-running services
 - [Offered by INDIGO Orchestrator + Mesos frameworks]
 - Run processing jobs
 - Provide DevOps approach
- Provide function as a service (Not fully available yet, possible future developments)
 - Execute code in response to events
 - Pre-processing triggered by new data availability (WIP in the INDIGO Orchestrator using storage events)

EOSC-hub Envisioned architecture

- Core services
 - API server
 - Workflow Engine
 - Message Bus
- Plugins
 - Cloud native interfaces
 - Container orchestration interfaces
 - HPC integration (e.g. QCG APIs)
 - Storage services integration
- Dependencies on Federation Services
 - EOSC-hub AAI
 - AppDB-IS (optional)
 - EOSC-Hub Monitoring (optional)
 - Marketplace (optional)





Integration with the Federation Services (1)

- AppDB-IS
 - Currently the INDIGO PaaS Orchestrator relies on the INDIGO CMBD service to retrieve information about the sites: service endpoints, available images, flavours, etc.
 - CMDBv2 under development (Cyfronet) is a fork of the AppDB-IS
 - Allows to extend the schema with extra information, e.g.
 - gpu, infiniband
 - public clouds
 - other compute services, e.g. mesos clusters
- **SLA** information can be exploited to support the automatic selection of the site for user deployments
 - No service available in EOSC-HUB yet
 - Currently the PaaS Orchestrator uses INDIGO SLAM. As soon as another solution is available, it will be plugged in.



Integration with the Federation Services (2)

- Monitoring information can be exploited to support the automatic selection of the site for user deployments
 - The PaaS Orchestrator consumes information coming from the INDIGO Monitoring pillar (based on Zabbix)
 - The monitoring metrics collected by the probes for the different sites and services are used to rank the providers and their services



• TOSCA (Topology and Orchestration Specification for Cloud Applications)

«an OASIS open standard that defines the interoperable description of services and applications hosted on the cloud and elsewhere; including their components, relationships, dependencies, requirements, and capabilities, thereby enabling **portability** and **automated management** across cloud providers regardless of underlying platform or infrastructure...»

- Designed to be extensible
- Adopted by the following
 - Products: Openstack Heat, Cloudify, Opera, Alien4Cloud, etc.
 - Projects: INDIGO-DC, SeaClouds, DICE, OpenTOSCA, etc.



- OAuth2.0
 - Authentication/Authorization through tokens
- **REST APIs**
 - Create a new deployment (described in TOSCA)
 - Get information about the deployments
 - Delete a deployment
- Json formatted request/response
- The availability of clients and bindings for some languages (python, Java, go...) can facilitate the integration with other tools (e.g. portals)



- Workload portability:
 - Can benefit from the adoption of the TOSCA standard
 - In some cases site specificities may require specific implementation → can be handled at the PaaS Orchestration level
- Workload extension:
 - Support the possibility to extend the running workload to another site
 - Use cases:
 - distributed multi-site elastic clusters
 - Cloud bursting: the application runs on a private cloud and bursts into Amazon (or another public cloud)

EOSC-hub Interoperability Guidelines

- **TOSCA** support for handling the deployment topology
- **OpenID-Connect/OAuth2.0** for managing AuthN/AuthZ
- Support for CMF APIs for connecting to the cloud providers
- Support for HPC (QCG, other?)
- Support for Data Services APIs
 - Which?
 - Process storage events?



- INDIGO PaaS
- Infrastructure Manager
- OpenStack Heat
- AWS CloudFormation
- Cloudify
- AWS Elastic Beanstalk

- OSCAR/SCAR
- AWS Lambda
- Apache OpenWisk



EOSC-hub Open points & Future activities

- AAI integration
 - The deployment workflow requires the access to different services and resources (compute/storage) on behalf of the user
 - Token exchange and delegation mechanisms work with single Authorization server
 - Main problem: the token validation in multi-proxy environment
- SLA Management and INDIGO SLAM evolution
- Integration with the Cloud Compute area
- Integration with portals, Science Gateways

Thank you for your attention!

Questions?



� eosc-hub.eu ♥ @EOSC_eu



This material by Parties of the EOSC-hub Consortium is licensed under a Creative Commons Attribution 4.0 International License.

Contact

marica.antonacci@ba.infn.it