EOSC-SYNERGY

EU Milestone: M2.5

Internal status update on progress of Integration Plan, Services Integration, Policies and Data integration

Document Identifier:	EOSC-SYNERGY-M2.5
Date:	31/08/2020
Activity:	WP2
Lead Partner:	KIT
Document Status:	APPROVED
Dissemination Level:	PUBLIC
Document Link:	

https://drive.google.com/file/d/1p2z2s1_3vqIP_SKuiA9HHEdvDEUq4MGN

Abstract:

This document is an EOSC-SYNERGY Milestone describing the update on the resource integration plan. The focus is made in HPC resource integration, and support to Container orchestration engines.

EOSC-SYNERGY receives funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 857647

I. Copyright Notice

Copyright Members of the EOSC-SYNERGY collaboration, 2019/2022.

II. Delivery Slip

	Name	Partner/Activity	Date
From	Marcus Hardt	KIT/WP2	28/07/2020
Reviewed by	Ignacio Blanquer Jorge Gomes	UPV/WP4 LIP/WP3	04/08/2020 04/08/2020
Approved by	PMB	ALL	28/08/2020

III. Document Log

Issue	Date	Comment	Author/Partner
v1	28/07/2020	Initial version	Marcus Hardt / KIT
v2	05/08/2020	Incorporated input	Marcus Hardt / KIT
v3	08/08/2020	Input	All Sites

IV. List of Acronyms

Acronym	Description
AAI	Authentication and Authorization Infrastructure
EOSC	European Open Science Cloud
HPC	High Performance Computing
HTC	High Throughput Computing
IM	Infrastructure Manager
TSes	Thematic Services
WP	Work Package

1. Milestone description

The purpose of this milestone is to refine and define the evolution of the infrastructure provided in WP2

- Integration Plan
- Services Integration
- Policies and Data integration

2. Current Status

The thematic services of WP4 have collected their requirements to the infrastructure in parallel to the initial integration of site infrastructures into EOSC in WP2. The result is included in this milestone as Table 1. The essential result of the table is that the infrastructure requirements can be categorised into a set of resource types:

1. HTC (grid) workloads do not have shared disk restrictions or special configurations. These will use the EGI HTC through the Workload Manager.
2. Virtual Machine workloads. Those will be instantiated through EGI Cloud Compute or specific OpenStack deployments, using the Infrastructure Manager (IM).
3. HPC workloads that require special configurations (such as Galaxy). These will instantiate a back-end in the cloud through IM and may use HPC resources via SLURM.
4. Container workloads. These require a management system such as Kubernetes as well as infrastructure resources, both can be instantiated via IM.

These requirements can all be met with the resources available in WP2. The above points 1. and 2. are already operational via the use of EGI Workload Manager (1.) and the EGI Cloud Compute and the integration of project members with the EGI AAI (2.). Planning for several details for HPC (3.) and Containers (4.) is still necessary, to meet the boundary conditions at the infrastructure operators, and to address the technical challenges that may arise.

3. Updated integration plan

This integration plan addresses services and data integration

3.1 Updated integration plan for HPC resources

HPC systems are expensive investments made by specific partners in WP2. Therefore, high requirements on the security of the authentication and authorisation systems are set in place. The new AAI that follows the AARC Blueprint Architecture is capable of complementing the multitude of existing local policies. Especially in the context of the project, the collaboration with WP4 ensures that policy requirements can be met. WP2 will focus on the technical challenges to enable the use of HPC resources.

The plan for the integration consists of a set of steps, working on which has already started:

- Evaluate and extend a way for enabling Secure Shell connections (SSH) with federated AAI to the HPC resources. (This is based on previous work inside the DEEP-Hybrid-DataCloud Project and current work.)
- Roll out solutions to participating sites.
- Leverage existing batch schedulers (SLURM in our case).
- Provide documentation to WP4.

3.2 Updated integration plan for Container Orchestration Engines

Kubernetes is not the only container orchestration system, but it is the one that most of the thematic services with this requirement asked for. While we as a project with a limited lifetime cannot sustainably influence the way in which partners are operating their infrastructures, we support the installation of partner's Kubernetes clusters.

To start gathering experience with providing Kubernetes to the Thematic Services, we will make use of **virtual Kubernetes Clusters** that will be instantiated with the IM and specific TOSCA templates.

The plan for this integration will consist of the following steps.

- Enable more users and sites to facilitate IM to instantiate virtual Kubernetes Clusters on top of the EGI Cloud Compute resource that are available in the project. (This is the current status.)
- Evaluate and extend the existing AAI Integration. Related work at the partner UPV has already developed a transparent integration with EGI Check-in, in which service accounts and namespaces are dynamically created through a reverse proxy and an authorizer module. This allows facilitating the VO concept for consistent authorisation.

- Discovery of virtual Kubernetes Clusters, to allow users and VOs easy discovery of resources that are available for them.
- Evaluation and integration of user interface for the cluster. The idea is to allow users to keep on using the kubectl CLI and work as if it were the local Kubernetes cloud.
- Extension of the solution to enable the efficient use of local storage. For this, we have run a survey on the available storage resources inside the project. (See Table 2)
- Extension of the solution to support elasticity of the virtual Kubernetes Clusters, so they can dynamically grow and shrink, depending on the workload.

4. Appendix

4.1 Service requirement table

The requirements of the thematic services have been collected by WP4 as part of their Deliverable D4.2

Service	WORSICA	G-Core	SAPS	Scipion	LAGO	SDS-WAS	UMSA	MSWSS	O3AS	OpenEBench
AAI	EGI Check in	Kerberos LDAP & CAS User/pwd	EGI Check in	EGI Check in	EGI Check in	B2ACCESS	EGI Check in & Life-science AAI	EGI Check in	EGI Check in	Life Sciences AAI
Workload Mng.	ArcCE, Batch (SLURM)	GCore+ K8s	K8s	Batch (SLURM)	Cluster Batch & EC3 / K8s / DIRAC4EGI	Batch (SLURM)	Batch (SLURM) in IM/EC3 (in Galaxy)	Batch (SLURM) in EC3 (in Galaxy)	Cluster batch (SLURM) & Docker Swarm / K8s	GA4GH WES/TES stack
Resource Mng.	IM (TOSCA)	IM / EC3	IM / EC3	IM / EC3	Cluster batch & IM+EC3	EC3	IM / EC3	IM / EC3	IM / Cloudify	ONE
Data Storage	Nextcloud, Dataverse	ElasticSearch for the catalogue	OpenStack Swift	Local	EGI DataHub ONEDATA	B2SHARE /B2SAFE	Local + S3	Local	WebDAV	Local + B2SHARE

Table 1: Services to manage each one of the four functionality blocks for each thematic service. Green denotes that they have been implemented at the moment of the deliverable, Blue boxes denote components that will be integrated by MS18. Grey boxes are functionality that have not been chosen yet. Underlined services are services listed in the EOSC marketplace.

From Table 1 we find the following ranked list of requirements:

1. SLURM (5 of 10): Direct HPC access. This is well known and depends on the HPC integration to finish
2. **Kubernetes (3-4 of 10): A container infrastructure for the TS/VO.**
3. Other (1-2 of 10): Galaxy? (OpenEBench) / Undecided (O3AS)

4.2 Storage availability table

The efficient integration of container resources with site-storage requires an adaptation of virtual Kubernetes clusters with real storage. These storages are available in at project partners:

Partner	Storage	Comment
BIFI	CEPH	
CETA-CIEMAT	NFS, CEPH	
CSIC	GPFS, CEPH, NFS, WebDAV	
CESGA	NFS, Lustre	
CESNET	CEPH	
IISAS	NFS, CEPH	
LIP/NCG	CEPH, Lustre, CVMFS	
PSNC	CEPH	
KIT	Webdav on OIDC	
CYFRONET	Onedata	To be confirmed

Table 2: List of available storages at partner sites