EOSC SYNERGY HANDBOOK

A Handbook targeted at Computer Centre Management, Administrators, and Users of the Infrastructure

Marcus Hardt and Viet Tran for the EOSC-Synergy WP2 Team

Table of contents

1. Introduction

In this handbook we describe how to integrate computing and storage infrastructure in such a way as to comply with current guidelines and best practices of the European Open Science Cloud (EOSC). We also describe how it may be used, operated and extended, with a focus as well on the system administrator's perspective. We present a set of example applications that were adapted to profit from the services supported by the infrastucture. This includes the services used, the general architecture of EOSC, the tools we chose to support the applications, and how computer centres can join their own resources into the federated cloud.

Our initial starting point was:

- We have 10 demanding Thematic Services that need infrastructure resource (cpu/gpu, storage, network, accounting, and monitoring) to provide their services and or results to the users
- We have an architecture for the European Open Science Cloud (EOSC) with definitions of the core services / federating core, ...
- We have an EOSC Marketplace that provides a large choice (320+) of different solutions and tools, many of which are in an unknown state.
- We have a team of experts in distributed computing that (happen to) operate the first working prototype of an EOSC infrastructure: EGI Federated Cloud

The challenge: Bring users and infrastructures together - in a scalable way that avoids vertical solutions. The EOSC-Synergy way to address this was the introduction of tools that provide a natural separation between the different roles and requirements on the infrastructure. Users are supported by the Community Manager, or by their Community Developers. These two representatives of the Community can request changes to the infrastructure services. Site administrators are responsible for the operation of the physical machines that provide CPU cycles and storage. They are in contact with Community Administrators that are in contact with Community managers to request the capacity in which these services are provided. Details are described in our Deliverable [D2.2].

All tools and solutions used are provided by open source software, exclusively. This ensures a long term perspective for a sustainable infrastructure. Furthermore, the open approach taken guarantees that custom extensions can always be implemented by third parties to create tools that bridge gaps that may be identified.

One example for this are the three different tools that may be used to access the infrastructure. Two different web-tools and one Command Line Interface (CLI) tool, address the whole bandwidth of user experience. These tools Infrastructure Manager (IM) and its dashboard, the Openstack Dashboard and the fedcloudclient are described together with many others in section 4.

Federated is not just distributed. The federated nature of the infrastructure brings several challenges that need to be addressed in order to build a sustainable solution. One particular challenge stems from the fact that our users are identified by entirely different legal entities in different continents. Making use of recently developed modern Identity and Access Management solutions (often called AAI) allows offering services to reliably identified users without the attached cost of user management.

Just as our users, are the computer centres that provide capacity to the cloud are federated across different legal entities in different countries, most of which are currently situated in Europe.

One essential concept for addressing this are Virtual Organisations (VOs). The VO Management is delegated to a user community. Community managers negotiate quotas for their VOs with individual resource centres. The infrastructure provides usage statistics (accounting) and monitoring on the granularity of the VO. VOs are probably the most cost efficient and scalable way of addressing federated users on federated infrastructures at large.

VOs as implemented on the EGI Federated Cloud provide a balance between the freedom in the authorisation decision and a strict governance on technology and policies (which software, which regulations, who is responsible). Without VOs, furthering endeavours such as the European Open Science Cloud EOSC do not seem feasible.

The overall organisation of this handbook is as follows. We start with an introduction of the EOSC-Synergy supported Thematic Services (TS) and their requirements. Each Thematic Service is described in section 2. Then we describe the general components of the EOSC architecture in Section 3.

The more technical Sections 4 and 5 describe how to integrate new resource centres (I.e. CPU, GPU, or storage hardware to provide cloud services) into the federated cloud, followed by a list of the components and tools used within the EOSC-Synergy Project. We close with the summary in Section 6.

2. Thematic Services

EOSC-Synergy is supporting ten different thematic services in four scientific areas (Earth Observation, Biomedicine, Astrophysics and Climate Change see here for a detailed list). Each service corresponds to a separate Community, each of which has different requirements, software tools, and access patterns. All of these thematic services require access to infrastructure services such as CPU, Storage and Network. The way in which the infrastructure is allocated, accessed, and used is different between each thematic service, though.

Many services provide access to research data and benefit from a wider availability. As such they need to act as clients and servers simultaneously. Often, the underlying user management (also called AAI) is shared for both roles.

The thematic services serve as examples, since they addressed a larger number of issues than many other services. This gives us the chance to either prove that EOSC-Synergy is ready to access federated datasets, in clusters distributed across Europe, or to develop additional tools for the ecosystem in case they are still missing.

During the project lifetime the communities have progressed towards best practices for the adoption of common EOSC guidelines, tools, interfaces and services. This includes strengthening the communities in increasing the capacity, performance, reliability and/or functionality of these thematic services through their integration in EOSC. This was especially important to increase the number of users of these thematic services substantially.

A detailed report will be published by EOSC-Synergy’s Work Package 4 (WP4), that describes requirements and solutions of the thematic services in detail. [WP4-IS]

2.1 Thematic Services Challenges

Here we provide a short overview about the specific challenges faced by each thematic service. These challenges regard access to Computing and Storage specifically. For more information about each specific thematic service, we provide references to a publication that describes the service in more detail. The following information have originally been collected in the paper “A survey of the European Open Science Cloud services for expanding the capacity and capabilities of multidisciplinary scientific applications” by Ignacio Blanquer et. al. [WP4-IS]

Thematic Service	Limitations and needs
WORSICA	- Improve download speed and number of concurrent downloads of satellite images.
	- Increase storage of the images needed for the algorithm.
	- Increase computational resources: GPU and RAM to speedup the image processing.
	- Seamless authentication and authorization for end users.
SAPS	- Need for a larger-scale deployment: computing, storage and data access.
	- Scalability and standardisation of services
	- Integrated and widely supported AAI
GCore	- Overcome limited access to data repository due to network bandwidth restrictions.
	- Infrastructure resources for processing and reprocessing large data sets.
	- Data delivery volume. Increasing size of files to be delivered to users.
SCIPION	- Insufficient Cloud resources for the workflow: GPUs, CPUs and RAM
	- Need of a Resource Management able to optimize the use of cloud resources.
	- Storage limitations and data transfer performance: 1-3 TB raw data.
	- Distributed and shared file system.
OpenEBench	- Need to work on heterogeneous systems to reach Life Sciences Communities
	- Need to efficiently store processed data and workflows in a FAIR manner.
LAGO	- Limitations on data preprocessing.
	- Needs data storage that copes with FAIR, curation and harvesting;
	- Need for computing power for simulations, together with optimal scheduling.
SDS-WAS	- Lack of services needed for Data storage and curation.
	- Lack of computing power for data analysis on-demand.
	- Lack of reliability of data sources, especially about observations
UMSA	- Long-term data storage is required, together with appropriate data curation.
	- Tracking provenance of the secondary (derived) datasets.
	- Need for reimplementing UMSA algorithms to deal with sparse data.
MSWSS	- Needs data protection measures because of the usage of confidential data.
	- The data has to be stored in a private storage only.
	- Implement security policies to protect VMs.
O3AS	- Requires larger storage resources, specially improving data availability
	- Fast handling of big data

Table 1: The Thematic services with their challenges that need to be addressed in EOSC Synergy

2.2 Thematic service technology choices

To address the identified challenges, WP4 of EOSC-Synergy undertook an analysis of the Services offered via the EOSC Marketplace. More than 320 services are available. In [WP4-IS] these services are organised into six categories, out of which "Access physical & eInfrastructures" is the one we are interested in. Table 2 shows those services chosen by each service to address the needs within the different categories.

More details are available in the corresponding WP4 Deliverable [D4.3].

Service	AAI	Workload Mng.	Resource Mng.	Data Storage
WORSICA	EGI Check in	ArcCE, Batch (SLURM)	IM (TOSCA)	Nextcloud, Datavers
G-Core	CAS User/pwd & EGI Check in	GCore+ K8s	IM / EC3	ElasticSearch
SAPS	EGI Check in	K8s	IM / EC3	OpenStack Swift
Scipion	EGI Check in	Batch (SLURM)	IM / EC3	Local + EGI DataHub
OpenEBench	Life Sciences AAI	WfExS + NextFlow	OpenNebula	Local + B2SHARE
LAGO	eduTEAMS + EGI Check-in	Batch (SLURM)	Local clusters + IM / EC3	EGI DataHub ONEDATA
SDS-WAS	B2ACCESS	Batch (SLURM)	Local clusters	B2HANDLE / B2SAFE
UMSA	EGI Check in & Life-science AAI	Batch (SLURM) in IM/ EC3 (in Galaxy)	IM / EC3	Local + S3
MSWSS	EGI Check in	Batch (SLURM) in EC3 (in Galaxy)	IM / EC3	Local + Dataverse
O3AS	EGI Check in	Batch (SLURM) & K8s	cluster	Local + WebDAV

Table 2: The solutions used by thematic services in the different domains. (from [D4.3])

3. Architecture of the EOSC

3.1 AAI

The integration with the Authentication and Authorisation Infrastructure (AAI) is the fundamental step, because this is the mechanism that delivers the identities (unique identifiers and group memberships) to the infrastructure. This is essential in order for the users to access the services in the first place. Furthermore, AAI is the logical “point” where users may be informed about the policies (like privacy notices, AUPs) and where access rights to services may be enforced.

In collaboration with EOSC-Hub and EGI we have recorded a procedure [AAI-Integration-Procedure]. There, services are first integrated with a demonstration instance of the EGI Check-in service for initial testing. Successfully tested services will then be moved to the production instance of EGI Check-in. This process is documented and supported in various procedures of EGI [EGI-Proc-09].

This integration will allow users to access all services, using their home-Identity, their community identities supported by EGI (six at the time of writing) or one of six social identity providers. In addition, Virtual Organisations (VOs) have been created [EGI-Proc-14] to support the thematic services and facilitate their resource allocation requests. For conducting trainings we envisage collaboration with the EGI training VO.

3.2 Monitoring

Every production service requires a monitoring system, so the desired efficiency and accountability can be verified at any time. After defining what should be monitored, it is easy to create a test and display its results. For EOSC-Synergy, Nagios core service will be used. One of the key aspects of Nagios is to get the information about a particular test.

With Nagios we can monitor various kinds of hosts, services and actions. From the simplest check on a web service (ping) to a more complex test on the service itself, such as testing outputs giving some input on APIs, websites using selenium or even testing if the login runs smoothly. This is all examples, since plugins can be written to all types of tests to check even the small functionality, using different programming languages.

In addition, contacts and alerts will be used to notify about problems detected by the Nagios service. This allows us to efficiently react to problems in a timely manner.

3.3 Accounting

The EOSC Accounting service collects, stores, aggregates, and displays usage information of HTC compute, storage space, cloud VM and data set resources. Resource Centres that are providing compute or storage to the EOSC infrastructure have to implement a collector (a stand-alone script or program, or a built-in function of their resource system), that gathers accounting metrics formatted into a standardised record format. These metrics are then transferred via a messaging service to the Accounting Repository, which stores and processes the data to produce aggregations that are then sent to the Accounting Portal for display.

The Accounting Portal retrieves topology information on how resource centres relate to national infrastructures and regions from the configuration management database (CMDB) and community affiliation from the AAI service to properly organise the accounting data. Information related to groups or VOs should also contain information about scientific disciplines to allow the portal to properly classify the resource usage. The Accounting Portal already is integrated with several other tools, such as GOCDB and REBUS for topology and geographical data, Check-In for the AAI, the Operations Portal for VO and scientific discipline information. It also uses X.509 certificates to map users to institutions and these to countries.

3.4 Information Provider

The information system collects data from the resource providers in a research infrastructure and makes it available for workload orchestration. In the EGI e-Infrastructure this is a fundamental service both for the HTC or Grid technology and for the EGI FedCloud. There are different solutions for each type of service within those technologies, e.g. both ARC and HTCondorCE have ad hoc provider implementations for gathering the information. For Federated Cloud, we use a unique implementation, coined as cloud-info-provider, to fetch data from the supported Cloud Management Frameworks (CMFs), notably OpenStack and OpenNebula. The cloud-info-provider component leverages the APIs exposed by those CMFs to get key information about the Cloud resource provider, such as the projects and images that any given VO is allowed to use.

Consequently, the data collected by the information providers is essential for the operation of the different workload management services available through EOSC, such as the INDIGO PaaS Orchestrator or DIRAC4EGI.

3.5 EOSC-Exchange / Marketplace

- 320+ Service, out of which the thematic service chose the most useful ones.
- Interoperability Framework
- Federation

4. Resources Integration (for managers)

This chapter addresses the aspects that need to be considered when a computer centre wants to join a part of their cloud resources with EOSC. After motivating why joining resources to the federated cloud, we describe the general access policies, the trust model, and provide technical pointers for technical staff to implement the actual integration.

The focus here is less on the hardware, but rather on the realisation that raw computing capacity is not really an asset without adequate means to access it and make it easily available to scientists. What delivers added value is the know-how about using and adapting these facilities to multiple tasks or scientific domains. Power is nothing without control.

In EOSC, the strong move towards harmonisation of infrastructures, (similar to Amazon AWS or Google GCS), leads to a unification of interfaces and access patterns. In EOSC these are tailored for the needs and the culture of scientific work.

Consequently, in EOSC-Synergy we adopted this trend of bringing users and providers together. The tasks of WP2 were to adapt the computing facilities to the requests of the Thematic Services. WP4 adapts the Thematic Services to the computing infrastructures.

4.1 Why should computer centres join the EOSC ecosystem

In the spirit of building an adequate infrastructure for science in Europe, we shortly present reasons for computer centres to join forces under the EOSC umbrella. Starting with a different perspective, of what would be the alternatives to joining a federated cloud.

One option is to buy the necessary capacity at commercial cloud vendors which is often a lot more expensive, especially, when it comes to longer term storage and data transfers. These points are important when FAIR data and OpenScience are taken seriously. Furthermore, the commercial support model does not include serving specialised requests which are sometimes required for significant advances. An endeavour like WLCG seems unrealistic to rely entirely on AWS, GCS or a Telekom Cloud.

Another option would be the continuation of the traditional model of providing custom solutions on site. This model ties users into custom solutions each of which is different from computer centre to computer centre. While this may be necessary for the optimal allocation of some HPC machines, the general drawback is that cost for both learning and supporting the individual solutions is too large in comparison to the EOSC offer.

On the contrary, there are many good reasons for computer centres to support EOSC.

Als already mentioned, an evolving ecosystem of tools that can readily be used with EOSC is available and growing. The general **open source nature of the tools**, and of the infrastructure, are one key element to ensure the long-term availability and extensibility of the ecosystem. The **same access pattern** can be used to allocate compute resources (CPU, Block/Object Storage, Network, Archive, ...) at any centre that is part of EOSC. The **synergies of this ecosystem** (unified user-support, software reusable across many places, identical policies and procedures) strongly outweigh the initial investment of time to join infrastructure resources with EOSC.

Use-cases that benefit from EOSC span the whole range from classic number-crunching including specialised hardware like GPUs often used for training Artificial Intelligence networks one one end of the spectrum, to web-servers that publish scientific results or that merely serve information pages.

An obvious benefit of the federated nature of EOSC is for example that critical services may be operated at different cloud sites for **increased fault-tolerance**. If these services reside in different centres in different countries or continents is merely a deployment detail.

Authorisation is definitely among the most important topic in the federated EOSC world. And of course computer centres that join EOSC will retain full control about who may use their resources or not. The authorisation model is based on the concept of Virtual Organisations (VOs). Access to resources is granted based on membership in a VO. The decision of which VOs are supported is with the provider of the resources. This corresponds to the concept of computing-proposals in HPC, where a successful proposal is allocated an amount of CPU-time. In most cases these proposals are assigned a group, which may consist of multiple members and is administered by the Principal Investigator (PI) of the computing-proposal. VOs work in exactly the same way, but they may be supported at multiple different EOSC sites at the same time. Details about the EOSC-AAI can be found in the "EOSC Authentication and Authorization Infrastructure (AAI)", published by the European Commission:

Monitoring and Accounting is another important tool within scientific computing. EOSC interfaces to several systems and collects them centrally. The collected data are published at the central Accounting Portal. The granularity is per VO to respect the privacy of individual users.

The **Unified set of policies** within EOSC organised and regulates the responsibilities of users, VO managers, site owners, etc. An international team of experts has collected policies from partners around the world, and structured them clearly, and in a flexible way in the so-called Policy Development Kit. These policies reflect the best of breed of what is used in production practice by many large infrastructures for decades. They are designed to provide fully GDPRS compliant templates. In addition, the flexibility allows individual computer centres to add specific clauses for users to accept, when using services from that site. Details about the technical policies will be given in subsection 4.2.2.

The following list summarises the benefits of joining EOSC

- Extended scalability (beyond the size of one local cloud), using the exact same EOSC compatible interfaces to access remote resources provided by other participants in EOSC.
- Extended availability for critical services that may be operated at different cloud centres in different countries.
- Access via a modern federated AAI (also called Identity and Access Management, IAM, in industry), offering stringent security at a low cost of operation.
- A large ecosystem of tools that is tailored to work on the federated EOSC infrastructure. Example: Creation of a dynamically scaling kubernetes cluster with only a few mouse clicks. (The full set of tools is described in section 5.)
- Included monitoring and accounting
- Efficient communication and user-support by clear separation between hardware providers, technology development and end-users.
- Professional service management procedures

4.2 Policies and concept for access to resources

4.2.1 National case studies

Case studies of the countries Czech Republic, Poland, Slovakia, Spain, The Netherlands, and the UK have been conducted as part of WP5 in EOSC-Synergy [D5.1, D5.2, D5.3]. The first document analyses the structure of each country regarding the stream of funding for research infrastructures, cloud computing sites are typically part of. The 2nd document provides recommendations to improve the uptake of EOSC and to extend the available infrastructure resources, while the last document analyses the impact of the recommendations given.

The bottomline of the recommendations given in [D5.2] regarding EOSC capacity extension are:

- Raising awareness of EOSC in general is necessary, including its Rules of Participation, related areas of Open Science such as PIDs, FAIR principles, etc.
- Creation and adoption of national policies for FAIR data should be supported
- Define roadmap/strategy and structural funding to guarantee stability/continuity of vital EOSC-related services
- Motivate the researchers, for example by adopting the system of giving credit to research by not only honouring the traditional publication, but also other scientific resources, including data and software
- Introduce a stable funding model independent from the projects
- Creation of the uniform governance model on the national level. Public governmental data services might be more integrated into the EOSC ecosystem
- Increase awareness of EOSC within the community
- Encourage national research agencies to contribute to the development of EOSC activities
- Disseminate successful results of EOSC applications, for example those virtual organisations dedicated to specific scientific disciplines, which can accelerate the adoption of EOSC infrastructure by national initiatives
- Increase awareness of EOSC in the scientific community
- Facilitate the registration of data and services, the allocation of resources to support the services, the verification of the quality and the adherence to the FAIR principles by providing tools, examples, tutorials and support teams.
- Leverage the highly distributed nature of the research infrastructures
- Make sure that the funding system for research sets apart enough structural funds for the continuity of Open Science support services
- Adapt the system of giving credit to research by not only honouring the traditional publication, but also other research outputs, including data and software

4.2.2 Access Policies

The policy on who can access the infrastructure with which share is entirely up to the respective owners of the hardware. Participating sites own the resources, hence they are in control. The technology used to enforce the authorisation decision is the Virtual Organisation (VO) model. VO managers are responsible for getting their VO authorised to use a given quota at every single site. As part of the agreement to support a given VO, a site may request specific agreements between the VO and the site. This may for example include specific Acceptable Use Policies (AUP) to be agreed upon by each user before being accepted as a VO member.

EGI defines three different types of access policies, that reflect the above:

- Policy-based access: Users are granted access based on policies defined by the EGI resource providers or by the EGI Foundation; such policies usually apply to resources being offered with “sponsored use” to meet some national or EU level objective; for instance, a country may offer resources with “sponsored use” to support national researchers involved in international collaborations.
- Wide access: Users can freely access scientific data and digital services provided by EGI resource providers.
- Market-driven: Users can negotiate a fee to access services either directly with EGI resource providers or indirectly with the EGI Foundation.

Within these definitions, services allowing access to rival resources (e.g. computing capacity or storage space) are usually provided under a policy-based or market-driven access policy. Services allowing access to non-rival resources (e.g. software packages or scientific data) are usually provided under a wide access policy. All access policies may not be available for each and every resource, service or scientific data set.

4.2.3 Technical Infrastructure policies

The legal relationships between the different stakeholders in the Infrastructure have evolved over the past 20 years to a stable and generally set of policies. These are referred to as “technical” policies in this document, because those were driven by IT security and other technical personnel. Yet, legal consultation took place, to ensure practical applicability of all policies and templates.

The policies are collected in the AARC Policy Development Kit. Note, that some policies are defined as “Policy Frameworks”. These frameworks merely define a list of criteria that need to be addressed by a given policy to conform to the framework. This allows the toleration of differences between policies in different countries or infrastructures. Table 3 provides an overview about the different policies, by whom they are defined, and to whom they apply.

Policy Frameworks

The following frameworks are considered best practice for Research Communities enabling federated access. They enable trust and promote attribute release from the wider identity federation.

- Sirtfi Trust Framework Sirtfi demonstrates that an organisation complies with baseline expectations for operational security and incident response in the context of identity federations. To mitigate risk, an Infrastructure may choose to restrict its interactions to only those federated organisations who are able to comply with the framework. As well as the Infrastructure itself supporting Sirtfi, it is highly recommended that each connected service supports Sirtfi.
- Research and Scholarship Entity Category Research and Scholarship identifies federated services that are operated for the purpose of supporting research and scholarship activity. Identity Providers demonstrate their support for research and scholarship by releasing a defined set of attributes for a user, including name, email address and additional low-risk information that may be useful for their activities [R&S]. It is recommended that entities adopt and use this category since many Identity Providers will not release user attributes to services that do not publish the Research and Scholarship Entity Category. REFEDS provide additional entity categories, such as “Personalized”, “Anonymous” and “Pseudonymous” to cater for additional use-cases and the related attribute requirements.
- GÉANT Data Protection Code of Conduct The Data protection Code of Conduct (DPCoCo) describes an approach to meet the requirements of the EU Data Protection Directive and (version 2) with the General Data Protection Regulation (GDPR) in federated identity management. The Data protection Code of Conduct defines behavioural rules for Service Providers which want to receive user attributes from the Identity Providers managed by the Home Organisations.

AARC Policy Development Kit

Here we give an overview over the policies contained in the policy kit, their meaning, purpose, and possible application.

The Policy Kit builds on the Snctfi framework [SNCTFI].

		Management	**Infrastructure Security Contact**	**User Community Management**	**Service Management**	**User**
Top Level	Infrastructure Policy	Defines & Abides by	Abides by	Abides by	Abides by	
Data Protection	Privacy Statement	Defines			Defines	Views
	Policy on the Processing of Personal Data	Defines	Abides by	Abides by	Abides by	
Membership Management	Community Membership Management Policy	Defines		Abides by		
	Acceptable Use Policy	Defines		Defines		Abides by
	Acceptable Authentication Assurance	Defines		Abides by	Abides by	
Operational Security	Incident Response Procedure	Defines	Abides by		Abides by	
	Service Operations Security Policy	Defines			Abodes by	

Table 3: Overview about the different policies, by whom they are defined, and to whom they apply.

- The top level Infrastrastructure Policy serves to bind the entire policy set and stipulates requirements on each of the participants; Management, Infrastructure Security Contact, User Community Management, Service Management (including the Proxy Operator) and the User. The top level policy identifies additional policy documents; in this case the five that are mandatory for Snctfi compliance. The Infrastructure may wish to define additional policies, such as Service Eligibility, Disaster Recovery, or Data Management; these policies should be linked into the Infrastructure Policy to ensure a coherent Policy set. Top Level Policy regulates the behaviour and activities of participants in the Infrastructure, and binds all other policies in a coherent whole. It explains the relevant terms, and instructs certain actions to be taken. The Infrastructure must have a Security Officer. All services must have a designated Security Contact. The communities must designate a Security Contact, and must ensure that all Community users will accept and abide by the relevant policies (which are all policies). This can be achieved, for example, by showing an Acceptable Use Policy (AUP) that contains links to all Infrastructure Policies. Naturally, this can technically be done by the Infrastructure’s services.
- Membership Management Policy is a set of rules for the Community on how User membership should be managed. The Community must define an AUP. The template is provided. The Community must properly manage their users’ membership life cycle, and must record all actions conducted on it. All the outlined actions must be followed (i.e. rules for Registration, Assignment of Attributes, Renewal, Suspension, Termination). The Community must take actions to ensure proper data protection and auditability.
- Acceptable Authentication Assurance Policy outlines the acceptable authentication assurance for the community, but also for the Infrastructure. The standard way of conveying this information is to use the REFEDS Assurance Framework (RAF]). The Community must define their own Assurance procedures, especially in relation to Identity Vetting. This may depend on the acceptable assurance levels demanded by services, e.g. services may request RAF Assurance Profile Cappuccino, and Community Manager must ensure that it is followed.
- Acceptable Use Policy defines conditions of usage of Infrastructure resources, but may additionally define rules for the Community itself. At the very least, Community must input their name and purpose. The Community may reuse the Infrastructure policy, if that is enough for them.
- Policy on the Processing of Personal Data outlines that proper measures must be taken to protect the personal data of users when using Infrastructure services, but it also instructs the Community to do the same. The Community must accept this policy, and must ensure that, if the Community has services integrated with the Infrastructure, must follow these rules.
- Privacy Policy Template is a template for all the services to use and follow.
- Incident Response Procedure is a set of rules to follow in case of a security incident. All Services must follow and abide by this procedure.

These policies and their templates can be found at the AARC Policy website. For EGI Federated Cloud the policies are linked here EGI-Policies. Additionally, there is a Moodle course that serves as an introduction for the PDK, and explains the purpose and usage of policies. The course allows one to organise and systematise the policy writing and implementation with the Infrastructure in order to properly manage users and properly provide services PDK-MOODLE. Everyone that needs to understand or create policies in federated research context is strongly encouraged to take the course. The course is also available as a YouTube playlist PDK-Playlisth

4.3 How to join the infrastructure

The Infrastructure is a complex setup of more and of less well connected services. Less connected services are typically HPC centres and specific Storage facilities at individual computer centres. Often, these are neither connected to the common AAI nor to any joint accounting system. In the spirit of this handbook we refer to the better connected services, where joining the infrastructure is a considerable amount of effort. The subsection is structured into one part for infrastructure providers and one for users.

4.3.1 Infrastructure providers: How to join as a computer centre

EGI Federated Cloud is a complex and well connected infrastructure. It follows the principles of major IT Service Management standards (FitSM), provides accounting, monitoring, Identity- and VO management, and more. A site willing to provide resources to the EGI Federated Cloud needs to be integrated with a variety of services, so that a minimum level of service quality can be guaranteed to the end-user. All necessary steps and procedures are documented in detail at the EGI Cloud Compute webpage. This integration can be grouped into three categories:

1. Organisational prerequisites:

- Join your national grid initiative (NGI) to obtain an entry in the GOCDB.
- Ensure you can support the relevant policies.

2. Technical prerequisites: Integration of cloud stacks into EGI FedCloud follows a well-defined path, with infrastructure services such as accounting, monitoring, authentication and authorisation, etc. These configurations make your site discoverable and usable by the communities you wish to support, and allow EGI to support you in operational and technical matters.

- Have a cluster of compute nodes available on which OpenStack will be installed.
- Install the required additional tools and services for monitoring, authentication, accounting, networking, ...

3. Allocation policies:

- Make decisions regarding which Virtual Organisations (VOs) you want to support. These decisions may be updated at any time.
- Direct your existing and/or local user communities to setup a Virtual Organisation.

4.3.2 Infrastructure users: How to join as a user or community

Allocation of resources (CPU/GPU hours, storage) is done at the Virtual Organisation (VO) Level. Users therefore must either be a member of an existing VO, or create one. Once a VO is supported at one or more sites, users can start using resources. While this is generally well described in the corresponding EGI Federated Cloud documentation. we give a general overview here. Also, the list of services and tools in section 5 will provide useful support to users at all levels.

The infrastructure provides multiple interfaces designed for different knowledge levels of users and for the different types of services. The cloud infrastructure provides access - you guessed it - to cloud resources. More specifically, these are OpenStack. resources, installed at multiple computer centres (distributed across Europe) called “sites”. One way to use these cloud instances is to use the OpenStack web interface (horizon) at every site. Since these are non-trivial to discover, the EOSC-Synergy dashboard. was developed for simpler discovery. The web interface offers access to all functionality of OpenStack, which includes computing, block- and object storage, and networking. Images available for running are provisioned via the AppDB. made available to the VO by a VO administrator.

More user-friendly access (everybody who knows the horizon web-interface knows there is room for improvement) is available via the Infrastructure Manager Dashboard. After initial configuration (documentation is available, including youtube videos. readthedocs. and moodle. users can easily deploy pre-configured VM infrastructures, including dynamic SLURM, hadoop or kubernetes clusters.

This may be understood as a starting point for the exploration of the infrastructure. All of it is available in a more technical way for automation via the commandline fedcloudclient. and REST interfaces, described in Section 5.

Persistent storage is available via the EGI Datahub. which allows multiple useful access patterns, including mounted filesystems and object storage.

This versatile cloud infrastructure may be used for a variety of use-cases. The spectrum of favourable patterns includes medium scale HPC (including GPU usage), on one end of the spectrum, via Portals that serve results of queries to large databases and conduct HPC analyses on request, all the way to traditional server hosting on the other end.

For any service in this environment, users will authenticate via EGI Check-In, to which they are redirected automatically. EGI Check-In offers to either authenticate via your home-organisation (e.g. the university you work at), or via a “Community-AAI” such as eduTEAMS, ORCID, GitHub, B2Access, Umbrella or Facebook. It is important to choose the correct one, because your VO membership information may come from the chosen community. What may be a bit confusing is that EGI is also a Community-AAI. To find your VO memberships, you need to choose different identities to log in (e.g. google at first, university later). To avoid confusion, it is important to remember the choices made.

Further features include features include:

- Global accounting that aggregates and allows visualisation of usage information across the whole federation.
- Monitoring of Availability and Reliability of the providers to ensure SLAs are met.
- Since the opening of the EGI Federated Cloud, the following usage models have emerged:
- Service hosting: the EGI Federated Cloud can be used to host any IT service as web servers, databases, etc. Cloud features, as elasticity, can help users to provide better performance and reliable services.

- Examples:
- NBIS Web Services,
- Peachnote analysis platform.

- Compute and data intensive applications: for those applications needing a considerable amount of resources in terms of computation and/or memory and/or intensive I/O. Ad-hoc computing environments can be created in the EGI cloud providers to satisfy extremely intensive HW resource requirements.
- Examples:
- VERCE platform,
- The Genetics of Salmonella Infections,
- The Chipster Platform.
- Datasets repository: the EGI Cloud can be used to store and manage large datasets exploiting the large amount of disk storage available in the Federation.
- Disposable and testing environments: environments for training or testing new developments.

- Example:
- Training infrastructure.

All these tools may be used and combined to develop individual solutions that may be tailored perfectly for each use case.

5. EOSC Synergy Services and Tools

Within EOSC Synergy, we have used, integrated and developed several tools. Some tools existed already, others were extended or developed from scratch. All of the tools have in common that they proved to be useful for the integration of the thematic services supported by the project. This chapter presents these useful tools.

Since all services are open source, their interfaces are open and may be used with several different tools. The tools presented here present a subset of possible solutions. They are not meant to be exclusive. All tools may benefit from combining them with others.

5.1 Overview

This table shows an overview about the tools presented in this handbook

	HPC	Cloud Compute	Storage	AAI	Q/A	Training	Generic
udocker	✓	✓					
SLURM	✓	✓					
Infrastructure Manager		✓	✓				
Fedcloud Client		✓	✓	✓			
Dynamic DNS	✓	✓					✓
EOSC Performance	✓	✓					✓
Cinder			✓				
Swift			✓				
RClone			✓				
EGI DataHub			✓				
B2Share			✓				
Core AAI/IAM				✓			
oidc-agent				✓			
mytoken				✓			
ssh-oidc				✓			
flaat				✓			
Vault				✓			
Pipeline as a Service					✓		
Quality Badges					✓		
Learn@Synergy						✓	
Online Training Platform						✓	
Video Conferencing Tool						✓	
Cloud Storage						✓	
Training Infra Mgmt						✓	
Jupyter Notebooks						✓	
Hackathon as a Service						✓	
Service Management	✓	✓	✓	✓	✓		✓
Monitoring	✓	✓	✓	✓			✓

	HPC	Cloud Compute	Storage	AAI	Q/A	Training	Generic
Accounting	✓	✓	✓	✓			✓

5.2 Services and tools for Cloud computation on EGI Federated Cloud

5.2.1 Infrastructure Manager

Infrastructure Manager is a tool that eases the access and the usability of cloud infrastructures by automating Virtual Machines Instances (VMI) selection, deployment, configuration, software installation, monitoring and update of Virtual Appliances. It supports APIs from a large number of virtual platforms, making user applications cloud-agnostic.

Infrastructure Manager is intensively used by Thematic services in EOSC-Synergy. The service was significantly improved during the project based on feedback from users. In addition, several new recipes for EOSC-Synergy services were developed and added to the dashboard.

Links:

- Infrastructure Manager Dashboard Service
- Service at EOSC Marketplace
- Documentation
- Training and Tutorial

5.2.2 Openstack Dashboard

OpenStack Horizon is a web-based graphical interface that users can access to manage OpenStack compute, storage and networking services. It allows service administrators to use this Dashboard to launch virtual machine instances, storage volumes or even manage their networks.

On top of this, EOSC-Synergy developed a dashboard that allows accessing all participating sites from one dashboard. The dashboard becomes the central web-based GUI interface for managing resources on all OpenStack sites in the project.

Links:

- EOSC-Synergy dashboard
- Documentation
- Training and Tutorial

5.2.3 FedCloud client

The FedCloud client is a command-line client designed for interacting with OpenStack services in the EGI infrastructure. The client can access various EGI services and perform many tasks for users. It includes managing access tokens, listing services, and command execution on OpenStack services located in the EGI Cloud infrastructure. The client was developed during the EOSC-Synergy project and has become the official client for EGI Cloud infrastructure.

FedCloud client is designed for using in shell scripts or Python programs. That enables sophisticated ways to automate tasks interfacing the cloud infrastructure. Complex tasks like listing all virtual machines owned by a user on all OpenStack sites in EGI Cloud infrastructure can be easily completed by simple scripts using FedCloud client.

Links:

- Package repository
- EOSC Marketplace
- Documentation
- Training and tutorial

5.2.4 Dynamic DNS

The Dynamic DNS service provides a dynamic Domain Name System (DNS) service for EGI Cloud infrastructure. Users can register their own meaningful and memorable host names, using a list of provided domains (e.g. fedcloud.eu, eosc-synergy.eu) and assign to public IPs of their servers hosted in EGI Federated Cloud. Simple login using EGI Check-in allows registering your own hostnames.

By using Dynamic DNS you can host services in EGI with meaningful service names and freely move their virtual machines (VMs) between sites without modifying configurations (federated approach). The hostnames also enable the services to get SSL certificates for improving security and privacy.

Some domains dedicated for EOSC-Synergy were added to the service for supporting Thematic services: o3as.fedcloud.eu, repository.fedcloud.eu, vm.fedcloud.eosc-synergy.eu, worsica.fedcloud.eosc-synergy.eu. The Dynamic DNS service is also integrated with Infrastructure Manager for deploying thematic services with registered hostnames from Dynamic DNS.

Links:

- Service
- EOSC Marketplace
- Documentation
- Training and tutorial

5.2.5 EOSC Performance

EOSC-Performance is a search-and-compare platform where you can upload and search through results from multiple benchmarks. By comparing the data acquired from benchmarks, you can evaluate and decide which computing infrastructure provider would give the best performance for your applications. The service is developed and supported by the EOSC-Synergy project.

For computing infrastructure providers, they can also submit new entries so users can find their services. The interface to the platform can be done through a web based Graphical User Interface (GUI) or through an API in case you want to automate or integrate the data with your project.

Links:

- Service GUI
- EOSC Marketplace
- Open-API self-documentation
- Training Video
- Documentation
- API Documentation

5.3 Services and tools for authentication

5.3.1 AAI / IAM Services

The cloud infrastructure relies on services that provide the Authentication and Authentication Infrastructure (AAI). More specifically, so-called "Community AAIs" as defined in the AARC Blueprint Architecture (BPA) are required for users to log into the EOSC services. This is fully in line with the EOSC Architecture. The EGI Federated Cloud uses EGI Checkin (https://aai.egi.eu) as its infrastructure proxy. This enables a large number of communities to use the Cloud. Within EOSC Synergy, we have successfully used the EGI, the GEANT eduTEAMS (https://eduteams.org), and the EUDAT B2access (https://b2access.eudat.eu) services for our users.

5.3.2 oidc-agent

oidc-agent is a set of tools to manage OpenID Connect tokens and make them easily usable from the command line. It follows the ssh design, so users can handle OIDC tokens in a similar way as they would do with ssh keys. If users are using or designing an API which relies on OIDC authentication like accessing OpenStack sites with the FedCloud client mentioned above, these tools will come really handy to them.

Links:

- Repository
- Documentation

5.3.3 mytoken

OIDC tokens are a very handy and secure way to handle user identification and authorisation between systems, especially in a federated environment such as EOSC. However, their short life is a problem on tasks where the execution time can be longer than the expiration time of the token. Such tasks are not rare in a scientific community such as EOSC. To solve these issues, mytoken was developed to provide OIDC Access Tokens for example to long-running compute jobs.

Mytoken is a web service to obtain OpenID Connect Access Tokens in an easy but secure way for extended periods of time and across multiple devices. Mytoken focuses on integration with the command line through a command line client but also offers a web interface for users who prefer managing their tokens with a browser. If you like oidc-agent and you need to execute long lived tasks on cloud or HPC, this tool is definitely for you.

Links:

- Service demo instance
- Documentation

5.3.4 ssh-oidc

ssh-oidc consists of a set of tools that allows (you guessed it) ssh with OIDC. This tool allows you to authenticate and log in to remote machines using your institution (or any other organisation) credentials instead of using a secret key or password.

Focused on usability, ssh-oidc is divided around several tools and libraries to mimic a subset of the popular ssh capabilities. The two main components are an SSH client wrapper: mccli designed to run on clients computers and the service motley-cue for mapping OIDC identities to local identities (to be run on the server where users are planned to log in).

Links:

- General repository
- mccli manual
- motley-cue manual
- Putty extension

5.3.5 Python API for AAI in services: flaat

Flaat is a simple python library that allows a straightforward implementation of REST interfaces that are well integrated with the AAI.

By using decorators, individual functions can be protected, so they may only be accessed by authorised users. Authorisation may be limited to VO and Group Membership as well as to the Assurance of a users identity.

Links:

- Github Repository
- Documentation

5.3.6 Vault Secrets Manager

Applications in EGI Infrastructure may need different secrets (credentials, tokens, passwords, etc.) during deployments and operations. The secrets are often stored as clear texts in configuration files or code repositories that expose security risks. Furthermore, the secrets stored in files are static and difficult to change/rotate. The secret management service for EGI Infrastructure is developed to solve the issues.

Links:

- Documentation

5.4 Services and tools for Cloud storages

5.4.1 OpenStack Cinder

If your cloud is hosted on an infrastructure managed by OpenStack, this type of storage will be the easiest you can access and use. It is available via OpenStack dashboard and will look just like a harddrive in your VM. However, note with this solution only the users and services with access to your VM can access the storage folder.

If you need to provide access to data to external users but you do not want to provide VM access, you probably have to look for another alternative. However, you can still use Cinder to extend your VM storage or combine it with other solutions providing the interface you like the most (e.g. Nextcloud).

5.4.2 OpenStack Swift

Another solution by OpenStack. If you need fast data access with infinite scalability (no need to reshape volumes) you probably should look into Object Storage technology. Swift (OpenStack) is the storage alternative to Cinder and probably one of your better options to work with object storage.

There are multiple ways to access Swift storage, however they might not be so intuitive. If you decide to use rclone to synchronise your filesystem with Swift, the tool EGI Swift Finder can implement all the discovery and configuration for you.

5.4.3 RClone

Rclone is a command line program to manage files on cloud storage. It is a feature rich alternative to cloud vendors' web storage interfaces. Over 40 cloud storage products support rclone including S3 object stores, business & consumer file storage services, as well as standard transfer protocols.

Rclone mounts any local, cloud or virtual filesystem as a disk on Windows, macOS, linux and FreeBSD, and also serves these over SFTP, HTTP, WebDAV, FTP and DLNA. Authentication and Authorisation will depend on the protocol you choose.

To facilitate the use of RClone, which was developed in a different context, a utility program "EGI Swift Finder" that sets up the environment for the EOSC context was developed. Links:

- RClone homepage
- EGI Swift Finder
- Documentation for RClone in the EGI context

5.4.4 Nextcloud

Nextcloud is a suite of client-server software for creating and using file hosting services. Software is free and open-source making anyone allowed to install and operate it on their own private server devices. Manage and access your files knowing your data is in your data centre, on a server managed by you or your team, rather than floating somewhere in the cloud. It is simple to install and deploy, for example in one of your hosts at the Federated Cloud.

Nextcloud is designed to be accessed via the web interface and WebDAV. Authentication via EGI Check In currently only works for the web interface. To use WebDAV and other protocols, currently passwords or OAuth2 Tokens have to be created in the web interface, before they can be used on the commandline.

Links:

- Nextcloud Homepage

5.4.5 EGI DataHub

A data management solution trying to provide High-performance with unified data access across globally distributed environments. If you have a very distributed cluster that your services need to access, this is probably an option for you.

The data organisation and sharing is similar to a filesystem, users organise their data in virtual volumes called spaces and share access between groups. To access your data you have multiple options such as web interface, CLI (command-line interface) or an API. Authentication and authorisation are based on OpenID Connect and SAML, supporting as well the usage of tokens at API level.

Links:

- EGI DataHub
- OneData Documentation

5.4.6 B2Share

B2SHARE is a user-friendly, reliable and trustworthy way for researchers, scientific communities and citizen scientists to store, publish and share research data in a FAIR way. B2SHARE is a solution that facilitates research data storage, guarantees long-term persistence of data and allows data, results or ideas to be shared worldwide. B2SHARE supports community domains with metadata extensions, access rules and publishing workflows. EUDAT offers communities and organisations customised instances and/or access to repositories supporting large datasets.

To manage your data there is a web interface and HTTP API. Authentication and authorisation are based on password or OIDC, using access tokens in the case of the API. Note that EUDAT encourages FAIR principles, so double check the privacy of your data (e.g. Metadata is always publicly available).

Links:

- B2Share technical documentation
- EUDAT service catalogue entry

5.5 Services and tools for HPC

5.5.1 udocker

udocker is a basic user tool to run simple docker containers in user space without requiring root privileges. It supports download and execution of docker containers by non-privileged users in Linux systems where docker is not available. It can be used to pull and execute docker containers in Linux batch systems and interactive clusters that are managed by other entities such as HPC and Grid infrastructures or externally managed batch or interactive systems.

udocker does not require any type of privileges nor the deployment of services by system administrators. It can be downloaded and executed entirely by the end user. The limited root functionality provided by some of the udocker execution modes is either simulated or provided via user namespaces. udocker is a wrapper around several tools and libraries to mimic a subset of the docker capabilities including pulling images and running containers with minimal functionality.

The performance of udocker beats - depending on the execution mode - most other container execution engines.

Links:

- Github repository
- Documentation
- ://www.youtube.com/watch?v=jEFiZghFkI

5.5.2 SLURM

The Slurm Workload Manager is a free and open-source job scheduler for Linux and Unix-like kernels, used by many of the world's supercomputers and computer clusters.

It provides three key functions:

- Allocate access to resources to users for specified durations of time so they can perform work
- Provide a framework for starting, executing, and monitoring work, typically a parallel job such as Message Passing Interface (MPI) on a set of allocated nodes, and
- Arbitrating contention for resources by managing a queue of pending jobs.

Slurm is the workload manager on about 60% of the TOP500 supercomputers. It was used by the thematic services to distribute their workloads -- either on HPC machines or on Clusters composed from Virtual Machines on the Federated Cloud.

Links:

- Slum webpage

5.6 Platform for Software and Services for Quality Assurance

The adoption of quality-based practices is one common challenge when it comes to developing software, especially in research environments. The SQAaaS platform provides researchers with a modular platform. It provides a variety of modules, all targeted at improving and assessing the quality of software.

The SQAaaS platform is based on abstract Quality Criteria for Software on the one hand (SQA) and on Services on the other hand (SerQA). Consequently, the SQAaaS platform implements the tools and pipelines that allow the verification of such criteria.

Here we briefly outline a subset of the modules of the SQAaaS platform. The full details are available in [D3.2] and [D3.4].

5.6.1 Pipeline as a Service

Pipeline as a Service is provided via the tool JePL (https://github.com/indigo-dc/jenkins-pipeline-library), which is a library to implement Software Quality Assurance (SQA) checks in Jenkins environments. It is meant to make it easier to configure the SQAaaS pipelines without knowing the Jenkins syntax. For this it provides a more simple solution to adopt a DevOps development practice by leveraging the YAML language to describe the criteria from the QA baselines to be assessed. A guided graphical process exists to create an initial configuration for a pipeline.

5.6.2 Quality Badges

The Quality Assessment and Awarding module analyses the compliance of a code repository with the quality baselines, and issues digital badges to certify if a minimum set of quality is achieved. Services and Software have different quality baselines, both of which are defined in [D3.4 section 5] Badges are issued as “gold”, “silver”, or “bronze”.

5.7 Services for online training

EOSC Synergy WP6 provides several tools dedicated to training.

5.7.1 Learn@Synergy

Learn@Synergy: a classic wordpress website for basic instructions and links to services, suc as catalogue of courses and training materials: https://learn.eosc-synergy.eu/

5.7.2 Online Training Platform

Online training platform: Based on the Moodle platform. It provides interactive courses with user forums to support community interactions among students and tutors as well as immediate feedback: https://moodle.learn.eosc-synergy.eu/

5.7.3 Video Conferencing Tool

Video Conference service to connect, talk or share the screen with other people. It only takes a minute to set up a new room and send invitations to the meeting: https://vc.learn.eosc-synergy.eu

5.7.4 Cloud Storage

Shared drive is a cloud space for users to securely store and synchronise files. The service is based on the open source NextCloud software and it is integrated with the AAI: https://drive.learn.eosc-synergy.eu

5.7.5 Training Infrastructure Management

Training Infrastructure Management service that allows self-deployment of cloud training infrastructure for a given training. It allows managing the virtual machines and accounts for training participants. The service is based on the Infrastructure Manager(IM) software, which deploys complex and customised virtual infrastructures on multiple back-ends.

5.7.6 Jupyter Notebooks

Jupyter Notebooks for Interactive computing which allows service developers to make use of interactive training technologies such as Jupyter notebooks.

5.7.7 Hackathon as a Service

Hackathon as a service (HaaS): is a platform that has been created within this project to facilitate the organisation of hackathons taking advantage of the EOSC infrastructure and accessible through the EOSC Portal. A hackathon is a sprint-like event in which computer programmers and others involved in software development (UI and graphic designers or project managers) collaborate intensively on software projects with the goal of creating a functioning product by the end of the event

More details on the online training services can be found in the corresponding deliverable of the Workpackage 6

5.8 Monitoring and Accounting Services

5.8.1 Service Management

Service management is traditionally done using the “Grid Operations Configuration Management Database (GOCDB)”. It provides a repository, portal and REST style API for managing Grid and Cloud topology objects such as sites, services, or downtimes. GOC is a central tool for IT service management, where all relevant information about participating computer centres is kept.

5.8.2 Monitoring

Site and service availability is monitored in the monitoring service ARGO. It deploys and runs checks against the infrastructure and collects information from low level items (hosts, services) to higher abstractions (groups, organisations). The monitoring data pass through an analytics engine to generate rich reports. The EOSC Synergy project created its own service level agreement with EGI, and is therefore present in ARGO as one group: https://argo.egi.eu/egi/dashboard/SLA/EGI_EOSCSYNERGY_SLA

5.8.3 Accounting

Accounting collects usage information of many different services inside the infrastructure. To confirm with privacy regulation, data is collected on the level of Virtual Organisations. It helps to assess how much storage, CPU hours, or Virtual Machines are used by any given Virtual organisation. Here we link to the use of EOSC Synergy resources throughout the project lifetime: https://accounting.egi.eu/