Future Generation Computer Systems 98 (2019) 171–179

Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

The EGI applications on Demand service

Gergely Sipos *, Giuseppe La Rocca, Diego Scardaci, Peter Solagna

EGI Foundation, Science Park 140, 1098 XG Amsterdam, The Netherlands

HIGHLIGHTS

- An open, online, extensible platform to provide computational applications.
- An X509 certificate based authentication system, called 'per-user subproxies'.
- Release as a production service in EGI for all science disciplines.
- Several application, science gateways, IaaS clouds and use cases are implemented.

ARTICLE INFO

Article history:
Received 15 June 2018
Received in revised form 26 February 2019
Accepted 6 March 2019
Available online 13 March 2019

Keywords:
e-infrastructure
Long-tail
SaaS
Cloud
HTC
EOSC

ABSTRACT

This paper describes the EGI 'Applications on Demand service', a new offering from EGI specifically for individual researchers, small research teams and early phase research infrastructures to support them in scientific data analysis. The described service is available through the EGI Marketplace and, through a lightweight registration and user identity vetting process, allows user-friendly access to a growing number of scientific applications, tools and application hosting frameworks (science gateways; Virtual Research Environments; portal) that are configured to use a dedicated pool of cloud computing and High Throughput Compute clusters donated by members of the EGI federation. The service operates as an open and extensible 'hub' for providers and user support teams who wish to federate and share applications and services with individual researchers or small, fragmented communities — typically called 'the long tail of science'. At the time of writing the service is under integration into the European Open Science Cloud through the EOSC-hub H2020 project.

1. Introduction

EGI [1] is an e-infrastructure collaboration that provides advanced computing and data services for researchers, trainers and innovators in academia and industry. The collaboration operates a publicly-funded e-infrastructure based on the federation of national e-infrastructures, currently comprising more than 300 resource centers from Europe and beyond. Members of the collaboration provide a well defined set of services via EGI:

1. Compute services: high-throughput computing (HTC) clusters, IaaS clouds and support for Docker containers to run compute- or data-intensive tasks and to host online services across institutional and national boundaries.
2. Storage and data services: services and facilities at over 300 locations that enable scientific communities to store, transfer, share, archive data within their community and with the broader public.
3. Training services: delivering training about EGI services, about IT Service Management (FitSM) and about the ISO27000 information security standard. Provisioning a federated cloud infrastructure for software training events.
4. Operational tools, protocols and processes to manage operational tasks across heterogeneous distributed infrastructures. This service is used by members of the EGI federation, and also by various international scientific communities who need to connect and share their members' resources and data in a coordinated manner.

EGI resource centers rely on the expertise of the 'EGI Foundation', a not-for-profit organization that coordinates EGI. The Foundation oversees areas such as infrastructure operations, user community support, contact with technology providers, strategy and policy development, flagship events and dissemination of news and achievements. The EGI collaboration has participants and associated participants drawn from representatives of national compute/data centers (so called NGIs) and Intergovernmental Research Organizations, such as CERN. These entities operate the physical resources and services that comprise the 'EGI infrastructure' and the 'EGI service catalogue.

* Corresponding author.
E-mail address: gergely.sipos@egi.eu (G. Sipos).

https://doi.org/10.1016/j.future.2019.03.009
0167-739X/

The EGI federation allocates resources (compute, storage) to scientific communities via 'Virtual Organizations' (VOs). A VO is the online representation of a scientific user group whose members are usually work in similar or related research areas, or are part of the same scientific collaboration, for which reason they need access to the same or similar applications, software services, datasets and underlying hardware capabilities. The biggest VOs of EGI represent experiments of the Large Hadron Collider (ALICE, ATLAS, CMS and LHCb) [2], the VIRGO experiment [3], the Cherenkov Telescope Array Observatory [4] or life science researchers from multiple countries and diverse background (biomed VO) [5].

Since its start in 2010, EGI has had well defined processes to allocate resources (i.e. create and operate VOs) for large, structured, international user communities. These communities have well established and long-term presence, and they are resourced well enough to sustain skilled IT support teams who can instantiate and operate VO services for the researchers. The most advanced research infrastructures from the ESFRI Research Infrastructure Roadmap [6] are the typical operators of the largest EGI VOs.

While serving international communities with long-term existence is still EGI's primary focus, over the years it became inevitable to integrate and offer services for individual researchers, small research teams and members of early phase research infrastructures who often struggled to access applications, compute and data services in their home countries or at their partner institutes, usually because their countries and institutes do not have sustained e-infrastructures. It was soon realized that the tools and resource allocation policies that were designed for long-living, structured communities are unsuitable for these type of users because they require users to:

- Obtain and use X.509 personal digital certificates from Certification Authorities (CAs) recognized by EGI resource providers.
- Join an existing VO that matches the requester's research subject/goals, or set up a new VO with providers that are willing support the user with compute and storage resources.
- Integrate relevant scientific applications with the VO resources and fulfil operational responsibilities, such as VO membership management, resource allocation negotiations and community/application specific service monitoring.

These individuals are often referred to as the 'long tail of science' [7] and they share the common characteristic of (1) missing dedicated arrangement to computational and storage resources and online applications and services to manage and analyse large amount of data, and (2) lacking the skills and experience of deploying and scaling applications to distributed computer architectures.

Over the past years we observed a fragmented and heterogeneous landscape of services operated by the NGIs for the long tail of science. Unfortunately many of these services have been dismissed because they became obsolete or because they pose a steep learning curve for non IT-experts [8]. Last but not least, there was a lack of coordination among the different national e-infrastructures with respect to the services they offer to the long tail: sometimes the same service/application was ported several times to national infrastructures, multiplying integration and operational costs.

For most of these users the management of X.509 certificates to access e-Infrastructures to run scientific applications is still considered a limiting factor, especially in those countries where Certification Authorities are not established. For both EGI and the resources provider stand-point, there is also a significant non negligible overhead to operate new VOs and to allocate the needed resources to target the scientific need of these users.

As a response to the unique needs of the long tail of science EGI designed and developed the 'Applications on Demand' service (also referred as "AoDs" in the rest of this paper) in the context of the EGI-Engage FP7 EU-funded project. In a nutshell, the EGI AoDs serves the large number of researchers and research teams across Europe who do not fit into any established "Virtual Organization" communities and do not have user friendly access to dedicated, community specific computational resources, storage and Virtual Research Environments (VRE) or science gateways.

From a technical point of view the EGI AoDs federates cloud and HTC resources, scientific gateways and application services from the NGIs and from other EGI partners, making these easily accessible for cross-national use via the EGI Marketplace [9] by any researcher. EGI AoDs operates as an open and extensible 'hub' for providers of e-infrastructure and scientific services who can easily share applications, application developer environments (VREs, science gateways) and storage and compute resources with individual users. For countries with high quality of national services, the EGI AoDs provides an additional channel to make services available. For NGIs that has no offering for the long tail of science the EGI AoDs can act as the primary e-infrastructure offering.

The service was designed in late 2014, demonstrated in November 2015, reached Alpha release for early adopters in January 2017, and was opened for the general public in April 2017. Since January 2018 the service is part of the service catalogue of the EOSC-hub H2020 project, which is a 33 million Euro initiative between 2018–2020 to establish the first building blocks of the future European Open Science Cloud.

This paper introduces the EGI AoDs, its enabling components and example use cases from real life experiences. Section 2 details the user and provider requirements that drove the system development. Section 3 describes the service architecture and design considerations. Section 4 provides details about the different system components that were developed/customized to establish the service. In Section 5 we describe the steps that applications hosting framework providers and resource providers have to follow to make their framework and/or the resource provider accessible within the EGI AoDs. In Section 6 we describe some of the real-life use cases that the service enabled since its opening. In Section 7 we compare the solution with related works. In Section 8 we present our plans for the extensions and improvement of the service, while in Section 9 we draw conclusions from the described work.

2. Requirements

The design of the AoDs began with collecting and documenting the characteristics and needs of the long-tail of science users, and the preferences (or sometimes constraints) articulated towards us by the national e-infrastructures. We draw this input from the NGIs who are directly supporting researchers in their own country and who provide services towards them. We complemented their input with experiences gained by the EGI.eu User Community Support team since 2010, mostly by engaging with users from countries without a local NGI. Our findings are summarized in this section and they served as requirements for our system:

1. Users' needs:

1.1. Zero-barrier access: Any user who carries out non-for-profit research should be able to get an account with a 'start-up' resource allocation to access scalable application services together the underlying cloud/cluster/storage resources.

1.2. 100% coverage: Anyone with Internet access should be able to become a user, without the need for personal travel to a CA to obtain special credentials (e.g. X.509 certificate).
1.3. Extendible and open: The service should be extendible with additional applications to support additional scientific disciplines and user groups. Such extensions should be possible via open interfaces and protocols.
1.4. User-centric: Support for users should be available in as many EGI member countries as possible. Training programme shall be run in interested NGIs to educate the local teams and turn them into user support nodes.

2. Service providers' preferences:

2.1. Realistic: Define an architecture that is implementable under the available effort levels (in the EGI-Engage project that supported the work). Reuse existing EGI technological building blocks as much as possible.
2.2. Secure: Provide as trustable user identity vetting, and user activity tracking protocols as possible (close to the existing solutions that are based on personal X.509 certificates).
2.3. Scalable: Be able to scale to 100 s of compute/data/application providers. (The size of the user group is limited by the number of providers and the size of their clusters, clouds, storages.)
2.4. Recognized: Have sufficient policies and tools that enforce the users to acknowledge the service in scientific publications resulted from the use.

3. Implementation

The high-level architecture of the AoDs is presented in Fig. 1. The front-end of the service is the EGI Marketplace [9] - a one-stop-shop where all the services of the EGI service catalogue can be browsed and accessed. AoDs is one of the services there.

The user has to register and create an account in the Marketplace at first login. Login is possible with academic institutional credentials (through EduGAIN), social media accounts (Google, Facebook, LinkedIn and ORCID) and with EGI SSO accounts [10]. Depending on the authentication mode some of the user's profile information are automatically retrieved by the Marketplace from the identity provider. Other attributes have to be filled manually by the user through a web form.

The user can submit an access request in the Marketplace to AoDs, specifying the application he/she wants to use and a short description of the use case (motivation). When a new service order request for the EGI AoD service is submitted through the Marketplace, the EGI staff member on shift assigns the request to one of the available Operators of the service — one from the user's local country, or one who has the most technical knowledge of the requested application. Routing the access request to the most relevant national or topical expert not only allows us to scale, but also to e.g. conduct conversation in local language and to connect long tail users to national support teams, nurturing sustainable relationships between researchers and national e-infrastructure providers.

The operator is tasked to:

- Verify user's profile identity and checks details of the request (duration, capacity, application, use case). The user's profile includes links to his/her departmental website, scientific publications, LinkedIn profile and these – together with the order details – helps the support team evaluate the order and identify spam or fake requests with a pretty high certainty. If needed, the support team can even contact the user in email or by phone to ensure the request came from a real researcher.

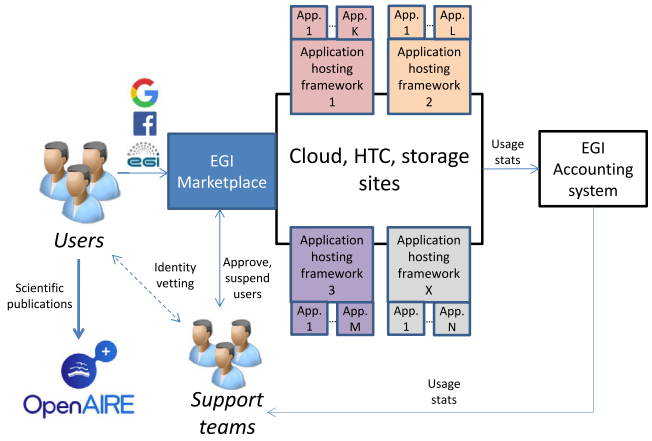

Fig. 1. The high-level overview of the EGI Applications on Demand service architecture.

- Respond to the service request, i.e. informs the user here and how to access the service.

Through a lightweight vetting process, the Operator will perform the needed checks to validate the user's request before to grant him/her the access to the large scale computing resources. In three working days, the service requests will be processed by the Operator and the user will be informed, via email, about the outcome of this evaluation. To each authenticated and authorized user is offered a grant, providing a pre-defined quota of resources, which can be used to run application(s) of choice. The quota to run applications, and user resources, can be renewed upon request. From a technical standpoint, the user's enrollment in the vo.access.egi.eu VO is done programmatically via the EGI AAI Check-In API. This registration process is completely transparent for the user's perspective. An unique and persistentUserId is also generated when the user logins one of the available Science Gateways and Applications hosting frameworks offered by the EGI AoDs service. The persistentUserId is used to generate a Per-User Sub-Proxy from a robot certificate (see more details in Section 4.4). Since 2017, more than 100 service requests have been processed by the Operators of the service.

Approved users are offered with access to the 'Application hosting environments' together with a default/custom quota of compute and storage resources on the connected 'cloud, HTC and storage sites' (See Fig. 1). The Application hosting environments are equipped with pre-configured applications, application developer/porting tools and appropriate user interfaces, manuals and user guides. For example, some gateways can be suited for workflow type applications, others for parameter study jobs, yet others for applications of certain scientific disciplines. The service is open for new application hosting frameworks and applications that providers think can be relevant for research communities (See Section 6 for more details on how to join). Science Gateways and Virtual Research Environments supported by the National Science Foundation and European Commission work programmes are ideal candidates for integration.

Approved users can login to any of the connected application hosting frameworks and use the embedded applications. These applications are already configured to scale to the distributed compute and storage systems. These compute and storage resources are either Infrastructure as a Service clouds; or are High Throughput Compute clusters, configured according to the EGI federational rules. Clouds shall follow the rules of the 'EGI Federated Cloud' [11], clusters should follow the rules of the EGI 'HTC Service'. The application use generates load on these resources

and this is reported into the EGI accounting system. From this accounting system the user support teams (as well as the users) can obtain statistics about individual users, as well as about application hosting frameworks, cloud/HTC/HPC sites. When a user exceeds the amount of compute/storage/network capacity that was allocated for him/her, the account can be suspended, blocking the user from further consumption through the service. To continue access the user has to submit a new justification which is again evaluated by the support team. (Or at this point the user may switch to a community-specific VO.)

Users of the service are asked (in the Acceptable Use Policy) to acknowledge the use of the service in scientific publications/presentations. Open Access publications are harvested by the web crawler service of the OpenAIRE scientific publication repository. Using the text mining algorithms of OpenAIRE, EGI (Support team in Fig. 1) is able to identify the papers that acknowledge this particular service so we can prepare statistics, reports for our contributors and funding agencies.

4. Enabling technologies

The service was created by customizing existing EGI components, and by developing a few new ones that glue together the elements into a single service. AoDs consists of the following components:

1. The EGI Marketplace
2. The EGI Check-In service
3. The pool resources of the EGI Infrastructure
4. The Per-User Sub-Proxy (PUSP) certificate
5. Scientific applications
6. Applications hosting frameworks (VREs, science gateways)
7. Policies

In our work we developed the PSUP component (#4 above), the Policies (#7 above), and integrated the other elements with these and based on these to result the overall service.

4.1. The EGI Marketplace

The 'EGI Marketplace' [9] was developed in the EGI-Engage project by Cyfronet (Poland) as a one-stop-shop where the catalogue of EGI-related services can be shared, ordered and accessed. The Marketplace is built on top of PrestaShop [12], a free and Open Source e-commerce solution. The Marketplace back-end is connected to a ticketing system through which service orders can be tracked and forwarded to providers for evaluation. Interested IT providers can propose new services for inclusion in the Marketplace through a web form. In the context of the EOSC-hub H2020 project [13], the same technology is used to pull together services from European e-Infrastructures, Research Infrastructures and other service provider communities with the ambition to become the landing page for users of the European Open Science Cloud.

4.2. The EGI Check-In service

The EGI Check-In service enables seamless access to EGI services and resources using federated authentication mechanisms. Check-In is operated as a connector entity for Identity Providers (IdPs) residing 'outside' the EGI ecosystem, and Service Providers (SPs) that can be either part of EGI, or associated partners. Check-In enables user login to the service providers with username–password. Check-In also aggregates user attributes from various authoritative sources (Identity providers and Attribute providers) and delivers them to service providers, helping them to make

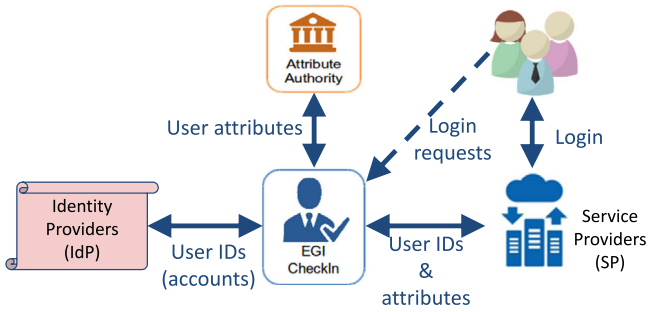

Fig. 2. The high-level architecture of the EGI AAI Check-In service.

informed authorization decisions about the authenticated users. The advantage of this design is that service and identity providers need to establish and maintain technical and trust relation only with a single entity, EGI Check-In, instead of maintaining many-to-many relationships.

In our scenario all the Applications hosting frameworks of AoDs are configured as service providers in Check-In, while Google, Facebook, LinkedIn, ORCID and EGI SSO are identity providers [10] (See Fig. 2). Check-In uses OpenID Connect (OIDC) or SAML2 protocol to interact with Service and Identity Providers, and uses the 'Per User Sub-proxy mechanism' (see later) to translate the username–password pairs to short-living X.509 proxy certificates that are used for authentication–authorization by the compute and storage sites.

4.3. The pool of resources from the EGI Infrastructure

A pool of cloud and HTC resources have been allocated to AoDs. At the time of writing the resource pool includes Infrastructure as a Service cloud resources from Italy (INFN-CATANIA-STACK) and Spain (CESGA) and HTC clusters from Belgium (VUB), Italy (INFN-Catania and INFN-Bari), Poland (CYFRONET) and Spain (CESGA). All the participating sites are joined together into an EGI Virtual Organization called 'vo.access.egi.eu' [14].

4.4. The Per-User Sub-Proxy (PUSP) certificate

The compute and storage resources are federated into the access.egi.eu resource pool with the use of the 10+ year old identity management system of EGI. This system requires short-living X.509 proxy certificates from the clients. The traditional method to generate such proxies is from a long-term personal 'Grid certificate', or from a long-term robot certificate [15]. While Grid certificates identify individual users, robot certificates identify applications that can be run on distributed resources by application portals on behalf of end users. Unfortunately neither of these approaches were sufficient for the Applications on Demand service, because:

- Personal Grid certificates proved difficult (sometimes impossible) to obtain and manage by users in various countries/regions and disciplines.
- Proxies generated from robot certificates do not include any information about the individual users, hiding all workload of an application under a single identity. This makes impossible to track individual users' activity (e.g. to recognize an excessive or harmful user).

To overcome these limitations we developed a new mechanism, the so-called 'Per-User Sub-Proxy' (PUSP) [16]. PSUPs are short-term proxies that are generated from robot certificates in a special way: The 'distinguished name' (DN) field of PSUPs includes

a unique string that is specific to the user for who the proxy was generated, and the DN is identical for any session of this user even if those sessions are in different 'Application hosting frameworks'. The compute and storage resources use the DN to authorize the user, and then to report resource consumption to the EGI accounting system. The user identity management and tracking process is graphically presented in Fig. 3.

The user-specific DN strings are generated during user approval and they are propagated to every participating application hosting framework when access is granted for a new user. The application hosting frameworks can generate PSUPs from robot certificates in one of the following two ways:

1. From a robot certificate that is deployed either locally on a USB smartcard on the server that hosts the framework.
2. From the 'eToken server' that runs at INFN-Catania. The application hosting frameworks can send a PSUP generation request to the server via its network API, including the userID as a parameter. The eToken server responds with the short-term proxy certificate that can be used by the framework to interact with the VO resources (see step 5 and 6 in Fig. 3). We setup this option because robot certificates are not available in every country but we wanted to be inclusive for application hosting frameworks. (Prior to our
3. work an application hosting framework must have had a robot certificate issued by the same country where the framework is operated. This limited the operation to countries where the national CA can issue robot certificates).

4.5. Applications hosting frameworks

Application hosting frameworks are hosted environments that provide user-friendly, application-specific or generic services for researchers. They often require the handling of 'big-data' therefore need big compute and storage resources. The frameworks use the identity federation of the Marketplace to enable access of approved users, and they use the PUSP mechanism to interact with cloud and HTC resources.

At the time of writing the AoDs includes three applications hosting frameworks: the WS-PGRADE/gUSE [17], the EC3/IM [18] and the CSG [19]. AoDs is open for any additional framework that wish to make applications and application development/hosting services available for a worldwide user base. Technical instructions to integrate a new applications hosting framework to the service are provided in Section 5.

4.6. Applications and tools

A total of 17 applications and tools from different scientific areas are already integrated into the applications hosting frameworks and are offered 'as services' to users:

- Molecular Docking, Workflow and parameter study tool (in the WS-PGRADE portal).
- Galaxy, Docker, Octave, Apache Tomcat, GnuPlot, NAMD, Hadoop, Marathon, Chronos, Jupyter Notebook, Cloud orchestrator (in the EC3/IM portal).
- Chipster, ClustalW2, Semantic Search, the Statistical R for Computing (in the Catania Science Gateway).

4.7. Policies

Two policies were developed:

- A security policy for resource centers that offer cloud/HTC/storage and application frameworks [20].
- An Acceptable Use Policy (AUP) [21] for users.

The first policy is compulsory to accept and implement by participating service providers. The policy defines that offering resources in this service shall not negatively affect the security other VOs supported by the centre. In particular, security incidents originating in AoDs should not impact the IT Infrastructure in ways that are incompatible with the operational model of other VOs. This document also provides guidelines on implementing the security procedures and controls at resource centers and in application hosting environments.

The Acceptable Use Policy (AUP) defines the conditions of use, and responsibilities of the users – such as using the services only for activities that relate to the work that was described in the access request form – and the text that should be used in scientific publications to acknowledge the service.

5. How to join as a provider

5.1. Application hosting frameworks (Gateways, Virtual Research Environments, etc.)

This section is a short summary of the main responsibilities of providers who wish to join AoDs. Full and up-to-date instructions are available at [22]. There are two fundamental prerequisites of integrating an application hosting framework in AoDs:

1. The framework must be a mature technology [23] with demonstrated use within publicly funded science. This requirement helps us keep EGI's 'production' status for research communities.
2. The framework must be able to use cloud, HTC or data services from EGI, or be ready to do so (to enable the scale-out of its applications).

If these are fulfilled then the provider should (1) register the framework as a Service Provider in Check-In using OIDC or SAML2 protocols.

- For SAML Service Provider, the authentication relies on the use of metadata. The metadata include information, such as the location of the service endpoints to invoke, and the certificates that will be used to sign SAML messages. The format of the exchanged metadata should be based on the XML-based SAML 2.0 specification [24]. Once the user is authenticated, EGI Check-in will return a SAML assertion to the application containing information about the authenticated user.
- For OpenID Connect Service Provider the EGI Check-in provides an OpenID Connect (OAuth2) API based on MITREid Connect [25] which has been certified by the OpenID Foundation [26]. The registration returns the OAuth 2.0 credentials (e.g.: clientID and secretKey) that will be used to authenticate users through the EGI AAI proxy. For applications hosting frameworks based on Liferay technology the INFN-Catania has developed an OpenID Connect module [27] that enables Liferay-based gateways to authenticate with OpenID Connect providers. This module is adopted by the WS-PGRADE/gUSE and CSG frameworks within this service.

As (2) step the framework should be made capable of generating user identifier proxies with DN according to the security policy. As it was discussed in the previous section the provider can do this either entirely within its own framework based on a local robot certificate, or by invoking the eToken server hosted at INFN-Catania.

If these two prerequisites are fulfilled then the provider with the help of the EGI team must:

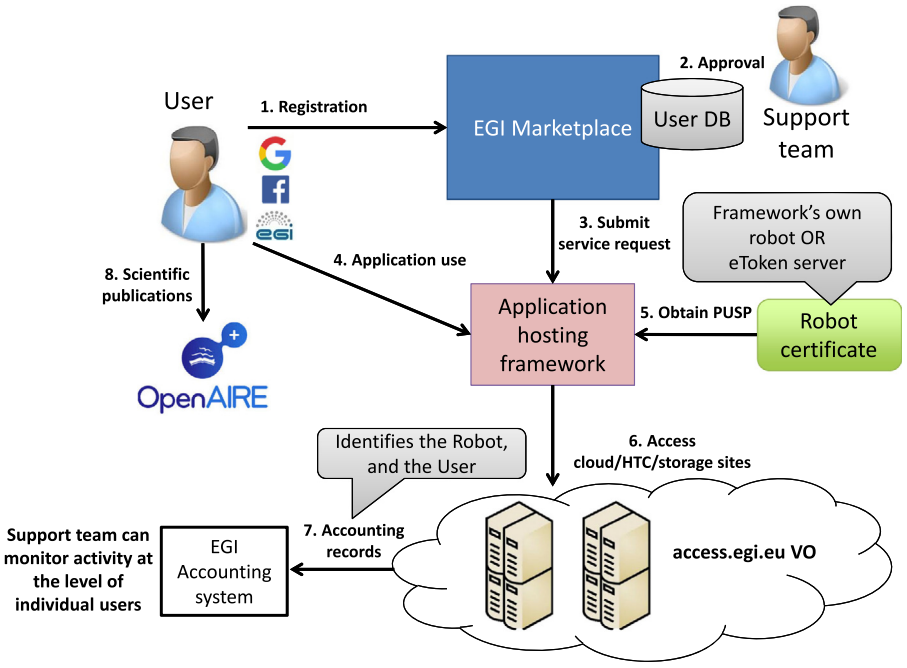

Fig. 3. Generating user identities, tracking user activities.

- Register the framework in the EGI service registry [28] to activate the availability and reliability monitoring for the framework with the EGI ARGO service [29].
- Setup a framework-specific support unit in EGI Helpdesk [30] (this is where tickets will be opened when the framework is noticed inaccessible by the monitor system).
- Sign an Operation Level Agreement (OLA) defining for example availability and reliability targets as well as helpdesk ticket response times.

5.2. Compute & storage providers

Compute and storage providers (cloud and HTC) can join the service by federating into the 'vo.access.egi.eu' Virtual Organization [14], following the regular EGI guides for resource providers. One unique extension compared to other Virtual Organizations is enabling the recognition of user-specific PUSPs. This requires a minor configuration in Openstack Keystone (Instructions for OpenStack providers) or OpenNebula rOCCI-server. Other types of EGI-compliant HTC services can do this 'out of the box'.

6. Success stories

Since the debut of the service in April 2017 we have received 57 service access requests from users located in 15 countries. These requests came from researchers working in Computer Sciences, Biological Sciences, Chemical Sciences, Mathematics, Social Sciences, Agronomy and Medical & Health Sciences. In the next subsection we briefly describe a few concrete use cases that profited from the service to demonstrate the diversity of requests attracted by the platform.

6.1. The recurrence quantification analysis (Massimo Rizzi, Mario Negri institute for pharmacological research, Milan)

Almost 40% of the annual incidence of epilepsy in humans originates from brain insults and the lack of an early marker of epileptogenesis (that is, a measurable event which can predict long in advance the development of epilepsy) impairs the development of novel therapies aimed at preventing the emergence of this neurological disorder. We have recently shown that during the early phase of development of epilepsy, long in advance of the occurrence of spontaneous seizures, the incidence of nonlinear dynamics in the EEG signals is increased [31,32]. To strengthen the possibility to exploit such dynamics as a general biomarker of epileptogenesis, we are currently exploring dimensional changes of nonlinear dynamics in EEG signals, with encouraging results. Our results provide evidence that dimensional changes during early epileptogenesis may be exploited to predict the ensuing epilepsy following the exposure to a risk factor and to develop novel anti-epileptogenic therapies (manuscript in preparation).

Results obtained considerably extend our previous findings [31,32] and EGI AoDs was crucial to this aim, since non-linear analysis requires significant computational power. The EGI AoDs provided us the computing resources to deploy virtual clusters on top of IaaS and use these clusters to perform non-linear analysis of EEG signals. During the several calculation sections the average number of vCPU cores was 40, so we could speed up our analytical algorithms to measure specific variables of nonlinear dynamics in EEG signals [33]. Indeed, the usage of VMs allowed reducing the time of processing of our data from about 3.5 years to less than 2 months.

6.2. Compounds for pain and inflammation treatment (Prof. Simone Ronsisvalle, department of drug science of the university of Catania, Italy)

Opioid analgesics are considered the main therapeutic agents for the treatment of moderate-to-severe pain. Acute pain is of sudden onset and is usually the result of a clearly defined cause such as an injury. Acute pain resolves with the healing of its underlying cause. Chronic pain persists for weeks or months and is usually associated with an underlying condition, such as arthritis. Inflammation is a response of the body triggered by a damage to tissues. In some circumstances the chronic inflammation can cause disruption of the tissues and it possible to provoke an excess of immune response. In all of these opioid receptor system

is involved, in particular, mu opioid receptor is the favourite target, and morphine is the preferred drug in the therapy of pain and in some cases for inflammation treatment.

In this scenario, the researchers synthesizes a series of benzomorphan amine substituted compounds. All of these compounds were computed for about 20 ns in the docking most stable and energetically favourable pose. The EC3/IM framework of the EGI AoDs has been used to deploy and configure an elastic virtual cluster with the NAMD application on top of. The cluster composed by 40 vCPU cores has been used to conduct molecular dynamic simulations. The outcomes collected by these simulations identified compounds that are very promising and have been taken to vivo/ex-vivo/vitro tests. Among them, one of the most promising compounds is in sub-nanomolar range of affinity for mu opioid receptor subclass with a very good selectivity towards delta and kappa. Additional molecular dynamic studies were conducted to evaluate the capability of the reference compounds to go out by the ligand binding domain. Other protein systems, including enzymatic systems, will be analysed using the grid. In order to evaluate the existence of new pathways with which the evaluated molecules can interact. This type of evaluation is very time consuming (10gg x 20 ns), but with the use of the grid this will shorten by about half. A scientific journal manuscript is in preparation.

6.3. Analysis of diffuse large B-cell lymphoma microenvironment (Maria Carmela Vegliante, hematology and cell therapy unit, IRCCS Istituto Tumori "Giovanni Paolo II" of Bari)

Diffuse large B cell lymphoma (DLBCL), the most common lymphoid malignancy in adults, is a heterogeneous disease with high variability in clinical outcome, genetic features and cells of origin [34–36]. A gene expression profiling study demonstrated that the "stromal-1" gene signature, enriched in genes encoding extracellular matrix proteins and macrophages, correlated with good outcome after chemo-immunotherapy [37]. To study the biological function of stromal cellular components in DLBCL, we set up a co-culture system by seeding together within the same 3D scaffold a fibroblastic-like and two different DLBCL cell lines. Upon a long-term culture, cells were retrieved and purified by magnetic beads carrying antibodies specific for each cell populations. We then performed gene expression profiling analysis by using Affimetrix platform (HG133plus_2) on RNA extracted from cell lines cultured alone and in combination with the another cell type.

The data were processed in the EGI AoDs, using the microarray analysis tools of the AoDs Chipster application. AoD represents a platform to support groups that lack expert bioinformatician and bioinformatic resources, as it includes user-friendly tools such as Chipster. The researcher got access to AoD by using her LinkedIn account and obtain permission to download and log-in Chipster just three working day after the formal request. AoD provided to the researcher detailed instructions to get access to Chipster testbed and supported her with efficient technical assistance during the following months of usage.

The analysis on Chipster testbed included a total of 16 samples organized in four groups (corresponding to the duplicates of four different conditions). Four CEL files (raw data) at a time were normalized using RMA normalization method and subsequently statistics tool were applied by using two group tests (default parameters). Interestingly, we obtained roughly 4000 genes differentially expressed between two different conditions that were visualized by heatmaps and principal component analyses (PCA). Data were exported as.txt file and used for further downstream analyses. Chipster was also used to analyse big public DLBCL datasets to perform PCA and to obtain results from GEP data of both small and big datasets.

6.4. Evaluation of a data-based interface for heterogeneous systems (Klaus-Dieter Schmartz, Germany)

Heterogeneity of data from different data sources presents a challenge for information derivation in big data applications. To address this challenge, the cloud compute resources of the EGI AoDs have contributed to scale up the configuration of a distributed streaming platform based on Apache Kafka. Using container virtualization techniques, the system is deployed in an Infrastructure-as-a-Service cloud. The user got access to the IaaS resource of AoDs and deployed Kafka through EC3/IM.

7. Related work

Science gateways have been used for more than a decade in distributed computing to host applications, to simplify access for users, and to offer higher/different abstractions for interacting with large-scale compute systems [38]. Because most of these gateways provide access to computational infrastructures that use X.509 grid certificates for user authentication, one of the recurring research topics was how to remove or hide X.509 certificates from the gateway layer that the users interact with. Several gateways simply ignore this problem and demand users to possess personal grid certificates [39,40]. Others introduced 'community accounts' [41] or robot certificates [42]. Both approach are based on the same concept: enforce strong authentication between the gateway provider and the compute resource, and make the gateway provider responsible for its users' actions. The gateway provider can make this arrangement only if it restricts the list of actions that the users can perform via the gateway on the compute resources. In practice this means the gateway becomes an application hosting environment, through which the users can run only the predefined list of applications. The gateway provider validates and keeps up to date the applications, therefore can lower the barrier of access towards the end users. Our approach enables users to still bring their own applications because we do not restrict their actions — rather we track the user actions with the use of per-user sub-proxies.

Compute services are operated for the long-tail of science in several NGIs of EGI, as well as in various non-European countries. These services are populated by local providers (i.e. providers from other countries cannot contribute) and are scoped for local users (i.e. usage from abroad is unavailable). EGI AoDs brings change in both aspects.

Services for cross-border access for the long tail are more prominent within the academic world for storage provisioning, particularly to deposit research data. Re3data.org is a catalogue of those data catalogues where researchers can upload and share data. At the time of writing re3data.org catalogues more than 2200 registries. Compute services are not catalogued alongside the registries.

The Analytics Environments on Demand (AEoD) from the University of California [43] is very similar to AoDS. AEoD operates high level applications for not IT-savvy researchers in an academic cloud. The service looks very much the same from the user perspective (high level applications that are accessible and scalable to multiple cores), there is a fundamental difference in the approach and service architecture: The Californian service is hosted and operated by a single institute, while the EGI service was designed to allow the federation of compute resources, gateways and applications from any number of distributed providers, integrating the 'best solutions' into an easy to access service.

A similar approach to AEoD, but from the commercial world is Alces Flight [44]. Alces Flight is a high performance computing (HPC) environment for research and scientific computing. Alces Flight includes over 1250 open-source applications and libraries

commonly used in research, as well as a cluster management system that to deploy scalable compute environments in compute clouds for those applications and libraries. The environment is compatible with major clouds, such as AWS. Alces Flight includes a much larger set of applications than AoDs, but it is maintained and provided by a single company as a downloadable software and as a service that can be started on-demand on third party commercial cloud providers.

A similar, but more data-driven and science discipline specific service is 'Copernicus Research and User Support (RUS) [45]. RUS is funded by the European Commission, managed by the European Space Agency, and operated by CS SI and its partners. RUS offers free access to Sentinel satellite data on an IaaS cloud through application toolkits that are offered in the form of virtual machines for the users. The toolkit includes various applications that are compatible with the Sentinel data and can be used online, within the cloud, eliminating the need for researchers to download and analyse data locally. Although RUS is offered by a small consortium, the cloud itself is single sited, while the user support is distributed (opposite to EGI AoDs where every element of the service is distributed).

8. Future work

The 'EGI Applications on Demand' service was opened for public use in April 2017. The promotion of the service to a continuous activity both by members of the EGI Foundation and by its partners within the NGIs. We are planning to update the technical architecture of the service in the coming years by replacing PUSPs with short-term proxy solution with the future RC Auth service [46]. RC Auth was designed by the AARC H2020 project [47] to have an open, European X.509 proxy factory service that can be used by any e-infrastructure and Research Infrastructure that needs user-specific X.509 proxies for service access. The use of RC Auth in the EGI Applications on Demand service would improve the compatibility of our architecture with other European initiatives, improving the sustainability and compatibility of our setup, by eliminating components that EGI currently sustains alone.

In the context of the HelixNebula-ScienceCloud H2020 project we recently integrated the service with the commercial IaaS clouds from T-Systems and Exoscale. The integration enables our users and our application providers to scale up applications even further, to the much larger IaaS clouds. During the HelixNebula-ScienceCloud project the extra capacity is available for users without any fees, and we will consider introducing access fees from mid-2019.

Between 2018–2020 the EGI Foundation coordinates EOSC-hub, a 33 million Euro H2020 project that involves over 100 institutes to establish the management system and the first services of the European Open Science Cloud (EOSC). Through this management system online and 'human' services, software and data will be delivered towards researchers via a single service Marketplace. The EOSC-hub consortium already includes providers from 3 e-infrastructure communities (EGI, EUDAT, INDIGO-DataCloud), and from 18 Research Infrastructures (mostly from the ESFRI Roadmap). EGI AoDs is one of the services in the EOSC-hub service catalogue and will be promoted via various EOSC-related channels in the coming years. In the EOSC context the AoDs will be promoted also as a platform that can be configured for a scientific community to bring together and make easily accessible specific gateways, scientific applications and cloud and HTC resources.

9. Conclusions

In this paper we presented a EGI Applications on Demand service. The service is specifically designed to cater for the needs of individual researchers, small research collaborations and early-phase Research Infrastructures. The service provides easy access to compute, storage and application resources without X.509 certificates, it is accessible from any country. The setup is open and extendible with new services through documented and open processes and policies. The system is scalable by connecting further IT resources and by scaling up the support team. The service eliminates the need for the long tail of users to form community-specific agreements with EGI providers.

The service does not replace the 'traditional', community/project specific Virtual Organizations. Large, structured scientific collaborations, Research Infrastructures and national e-infrastructures continue to require dedicated VOs because only those can provide fully customized, community-specific, dedicated services and compute/storage capacity for users. The EGI Foundation therefore continues to negotiate and secure services for community-specific VO through Service Level and Operational Level Agreements (SLAs, OLAs). The Applications On Demand Service is now part of the EOSC-hub service catalogue and will be promoted to EOSC users in the coming years.

CRediT authorship contribution statement

Gergely Sipos: Conceptualisation, Funding acquisition, Methodology, Supervision. **Giuseppe La Rocca:** Investigation, Software, Project administration, Validation. **Diego Scardaci:** Investigation, Software. **Peter Solagna:** Investigation, Software, Supervision.

Acknowledgements

The work was supported by the EGI-Engage H2020 project (grant number: 654142) and the EOSC-hub H2020 project (grant number: 777536). The authors would like to thank the contribution to the implementation of the described concept to the following institutes: CESGA, CYFRONET, INFN-Catania, INFN-Bari, NIKHEF, MTA SZTAKI, ULB-VUB and UPV. We also want to extend our warm thanks to Massimo Rizzi from the Department of Neuroscience of the Mario Negri Institute of Milan, Simone Ronsisvalle a Ph.D. Medical Chemistry researcher at the University of Catania, Maria Carmela Vegliante, a senior researcher at the IRCCS Istituto Tumori "Giovanni Paolo II" of Bari, and Klaus-Dieter Schmartz for describing their use cases.

References

[1] EGI: Advanced Computing for Research. https://www.egi.eu/, 2018 (accessed 27 2018).
[2] Daniele Bonacorsi, et al., WLCG Service Challenges and Tiered architecture in the LHC era, 2007, http://dx.doi.org/10.1007/978-88-470-0530-3_68.
[3] B. Caron, et al., - The VIRGO interferometer - in Classical and Quantum Gravity 14 (6) (1999) 1461, http://dx.doi.org/10.1088/0264-9381/14/6/011.
[4] L. Ambrogi, et al., On the potential of Cherenkov Telescope Arrays and KM3 Neutrino Telescopes for the detection of extended sources, 2018, http://dx.doi.org/10.1016/j.astropartphys.2018.03.001.
[5] T. Glatard, et al., Multi-modality image simulation with the virtual imaging platform: Illustration on cardiac MRI and echography, in: IEEE International Symposium on Biomedical Imaging (ISBI), Barcelona, Spain, 2012.
[6] European Strategy Forum on Research Infrastructures (ESFRI). Infrastructure roadmap 2018. http://www.esfri.eu/roadmap-2018, 2018 (accessed: 27 2018).
[7] Unwinding the 'Long Tail' of Science. https://www.ci.uchicago.edu/blog/unwinding-long-tail-science, 2018 (accessed: 27 2018).
[8] The EGI Accounting data about 'catch-all' Virtual Organisations. https://accounting.egi.eu/discipline/Miscellaneous/normcpu/VO/DATE/2014/5/2018/4/, 2018 (accessed 27 2018).
[9] The EGI Marketplace. https://marketplace.egi.eu/, 2018 (accessed 27 2018).

[10] The EGI SSO. http://egi.eu/sso, 2018 (accessed 27 2018).
[11] The EGI Cloud Compute service. https://www.egi.eu/services/cloud-compute, 2018 (accessed 27 2018).
[12] PrestaShop – free e-Commerce software. https://www.prestashop.com/, 2018 (accessed 27 2018).
[13] The EOSC-hub project. https://eosc-hub.eu/, 2018 (accessed 27 2018).
[14] The vo.access.egi.eu Virtual Organisation in the EGI Operations Portal: https://operations-portal.egi.eu/vo/view/voname/vo.access.egi.eu, 2018 (accessed 27 2018).
[15] IGTF Robot certificates. https://www.eugridpma.org/guidelines/robot/, 2018 (accessed 27 2018).
[16] The Per-User sub-Proxy. https://wiki.egi.eu/wiki/Usage_of_the_per_user_sub_proxy_in_EGI, 2018 (accessed 27 2018).
[17] P. Kacsuk, Z. Farkas, M. Kozlovszky, G. Hermann, Á. Balasko, K. Karóczkai, I. Márton, WS-PGRADE/gUSE generic DCI gateway framework for a large variety of user communities, J. Grid Comput. 10 (4) (2012) 601–630.
[18] M. Caballer, D. Segrelles, G. Moltó, I. Blanquer, A platform to deploy customized scientific virtual infrastructures on the cloud, Concurr. Comput.: Pract. Exper. 27 (16) (2015) 4318–4329.
[19] V. Ardizzone, R. Barbera, A. Calanducci, M. Fargetta, E. Ingrà, I. Porro, G. La Rocca, S. Monforte, R. Ricceri, R. Rotondo, D. Scardaci, A. Schenone, The DECIDE science gateway, J. Grid Comput. 10 (4) (2002) 689–707.
[20] The EGI Applications on Demand (AoD) Security Policies. https://documents.egi.eu/public/ShowDocument?docid=2734, 2018 (accessed 27 2018).
[21] Acceptable Use Policy and conditions of use of the EGI Applications on Demand service. https://documents.egi.eu/public/ShowDocument?docid=2635, 2018 (accessed 27 2018).
[22] Applications on Demand service – Information for providers. https://wiki.egi.eu/wiki/Applications_on_Demand_Service_-_information_for_providers, 2018 (accessed 27 2018).
[23] Technology Readiness Levels by the European Commission. https://ec.europa.eu/research/participants/data/ref/h2020/wp/2014_2015/annexes/h2020-wp1415-annex-g-trl_en.pdf, 2018 (accessed 27 2018).
[24] Metadata for the OASIS Security Assertion Markup Language (SAML) v2.0. https://docs.oasis-open.org/security/saml/v2.0/saml-metadata-2.0-os.pdf, 2018 (accessed 27 2018).
[25] MITREid Connect. https://github.com/mitreid-connect/OpenID-Connect-Java-Spring-Server, 2018 (accessed 272018).
[26] OpenID Certification. http://openid.net/certification/, 2018 (accessed 27 2018).
[27] The OpenID Connect for Liferay. https://github.com/csgf/OpenIdConnectLiferay, 2018 (accessed 27 2018).
[28] The GOCDB. https://goc.egi.eu/, 2018 (accessed 27 2018).
[29] The ARGO. http://argo.egi.eu/, 2018 (accessed 27 2018).
[30] The EGI Helpdesk. http://helpdesk.egi.eu/, 2018 (accessed 27 2018).
[31] M. Rizzi, F. Frigerio, V. Iori, The early phases of epileptogenesis induced by status epilepticus are characterized by persistent dynamical regime of intermittency type, in: C.L. Webber Jr., et al. (Eds.), Recurrence Plots and Their Quantifications: Expanding Horizons, Vol. 180, Springer Proceedings Phys, 2016, http://dx.doi.org/10.1007/978-3-319-29922-8_10.
[32] M. Rizzi, et al., Following a potential epileptogenic insult, prolonged high rates of nonlinear dynamical regimes of intermittency type is the hallmark of epileptogenesis, Sci. Rep. 6 (2016) 31129, http://dx.doi.org/10.1038/srep31129.
[33] The Mario Negri – Institute for Pharmacological Research, Milan. http://www.marionegri.it/en_US/home, 2018 (accessed 27 2018).
[34] Alizadeh, et al., Nature 403 (2000) 503–511.
[35] Monti, et al., Cancer Cell 22 (2012) 359–372.
[36] Morin, et al., Blood 122 (2013) 1256–1265.
[37] Lenz, et al., N. Engl. J. Med. 359 (2008) 2313–2323.
[38] N. Wilkins-Diehr, Science gateways—common community interfaces to grid resources, Concurr. Comput.: Pract. Exper. 19 (6) (2007).
[39] S. Gesing, J. Krüger, R. Grunzke, L. de la Garza, S. Herres-Pawlis, A. Hoffmann, Molecular simulation grid (MosGrid): A science gateway tailored to the molecular simulation community, in: P. Kacsuk (Ed.), Science Gateways for Distributed Computing Infrastructures, Springer, 2014.
[40] V. Korenkov, I. Pelevanyuk, P. Zrelov, A. Tsaregorodtsev, Accessing distributed computing resources by scientific communities using DIRAC services, DAMDID/RCDL, 2016, pp. 110-115.
[41] Von Welch Jim Barlow James Basney Doru Marcusiu Nancy Wilkins-Diehr, A AAAA model to support science gateways with community accounts, in science gateways—common community interfaces to grid resources, Concurr. Comput.: Pract. Exper. 19 (6) (2007).
[42] V. Ardizzone, et al., Science gateways for semantic-web-based life science applications, in: Studies in Health Technology and Informatics, Volume 175: HealthGrid Applications and Technologies Meet Science Gateways for Life Sciences, 2012, pp. 129–130.
[43] J.C. Christopher, Analytics environments on demand: Providing interactive and scalable research computing with windows, in: Proceedings of the Practice and Experience in Advanced Research Computing 2017 on Sustainability, Success and Impact, ACM, 2017.
[44] Alces Flight Compute. https://alces-flight.com/, 2018. (accessed 13 2018).
[45] Copernicus Research and User Support (RUS) Service: https://rus-copernicus.eu/ (accessed 13 2018).
[46] RCauth: The white-label Research and Collaboration Authentication CA Service for Europe. https://www.rcauth.eu/, 2018 (accessed 27 2018).
[47] The AARC H2020 project. https://aarc-project.eu/, 2018 (accessed 27 2018).

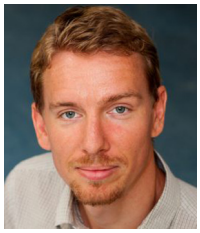

Gergely Sipos works as Customer and Technical Outreach Manager for EGI.eu. He manages EGI's user engagement, training and user support programmes, enabling scientific communities in exploiting EGI services to push scientific boundaries. He holds M.Sc. and Ph.D. in computer science with specialization in management, from the University of Miskolc (Hungary).

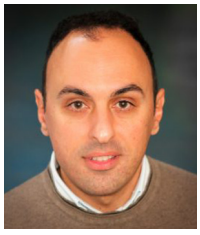

Giuseppe La Rocca works as Technical Outreach Expert at the EGI Foundation since December 2016. One of his main activities is to assist EGI in identifying and supporting research communities with their scientific applications from various domains. He is the Chief Community Officer of the EGI.eu User Community Board (UCB) group, and the EGI Applications on Demand (AoD) service owner. Formerly he worked as technologist for the Italian National Institute of Nuclear Physics (INFN) division of Catania. Since 2004, both at National and European level, he has been involved in several projects co-funded by the European Commission. During these years, he has matured strong skills and competences on Grids and Clouds technologies, working on ICT scientific developments for supporting both emerging and already established VRCs. Giuseppe holds a M.Sc. in Computer Science Engineering from the University of Catania (Italy).

Diego Scardaci works as Senior User Community Support and Outreach Officer at EGI Foundation. Diego has been involved in several projects, related to Research Infrastructure and e-Infrastructure and co-funded by the European Commission in the context of FP6, FP7 and Horizon2020, fulfilling several management roles since 2006. He is currently co-chair of the EOSC-hub activity management board. His expertise covers many areas: user support, operations and services development. Before joining EGI, Diego worked with INFN and in the R&D center of the Telecom Italia group. Diego holds a Master's degree in Computer Science Engineering from the University of Catania.

Peter Solagna has a master degree in Computer Engineering from the University of Padua (Italy). He has got several years of experience in software development and operations management applied to high performance and distributed computing infrastructures. Peter worked for the EGI foundation from 2011 to 2017 leading the operational activities of the EGI production infrastructure.