Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 68 (2015) 196 – 205

Procedia Computer Science

HOLACONF - Cloud Forward: From Distributed to Complete Computing,

The EGI Federated Cloud e-Infrastructure

Enol Fernández-del-Castillo[a,b,*], Diego Scardaci[a,c], Álvaro López García[b]

[a] *EGI.eu, Address, Amsterdam, Netherlands*
[b] *IFCA (CSIC-UC), Santander, Spain*
[c] *INFN Division of Catania, Catania, Italy*

Abstract

The EGI Federated Cloud is a standards-based open cloud system that offers a scalable and flexible e-infrastructure to the European research community, and extends the EGI computational offer beyond the traditional High Throughput Computing of the grid platform with new models like long-lived services and on demand computation. The EGI Cloud technology enables the federation of institutional clouds to run scientists workloads, simulations and services spanned across multiple administrative locations, allowing the users to access and exploit its resources as a unique system. The architecture of the federation was defined after a two year period of development based on a set of user stories describing operations on a cloud infrascturues and was officially launched in May 2014. Since then the EGI Federated Cloud operates as a federation of heterogeneous Infrastructure as a Service type clouds with every participating provider implementing the same set of interfaces towards users and system administrators. Enforcing cloud technology agnosticism and of supporting service mobility by means of open standards has also allowed for the inclusion of commercial cloud providers into an infrastructure previously supported only by academic institutions. This contributes to a wider goal of the funders to create economic and social impact from supported research activities. This paper details the state of the art, the design and development processes, and the organisational effort that have lead to the creation and deployment of the EGI Cloud Platform.

Peer-review under responsibility of Institute of Communication and Computer Systems.

Keywords: Cloud Computing; e-infrastructure; Federation of Resources; support

1. Introduction

Modern scientific challenges require cross-crounty collaborations and computing power with flexible usage to analyse vast amounts of data. EGI operates one of the largest, collaborative e-infrastructures in the world that allow scientists to share information securely, analyse data efficiently and collaborate with colleagues worldwide. EGI provides a uniform use oriented and collaborative access to computing and and data storage resource in mode than 30 countries, from EGI's National Grid Infrastructures (NGIs).

The EGI community started developing a new type of infrastructure, the EGI Federated Cloud in 2011. Following the successful use of the grid computing platform for the LHC Computing activities [1] and other communities requiring

* Corresponding author. Tel.: +34 942206779
E-mail address: enol.fernandez@egi.eu

Peer-review under responsibility of Institute of Communication and Computer Systems.
doi:10.1016/j.procs.2015.09.235

batch-oriented High Throughput Computing, EGI investigated how to broaden the support for different research communities and their application design models by enriching the solutions being offered, whilst retaining and protecting existing functionality and investment made in EGI's production infrastructure. The Infrastructure as a Service (IaaS) cloud service model was considered as clear candidate to enable this transformation.

The development of the EGI Federated Cloud was started by an EGI Taskforce, created to evaluate the integration of virtualised resources within the existing EGI production infrastructure. The taskforce started with the provision of a testbed open to all research user communities and cataloguing the requirements for a cloud federation. Two main design choices were made: platform agnosticism and open standards. The former allowed for retaining the expertise developed by many EGI providers in managing and deploying local cloud-based services, the latter to build on a number of on-going European and global activities that were engineering diverse open source Cloud Management Frameworks (CMF) [2,3,4]. As a consequence, every cloud provider of the EGI infrastructure was not mandated to use a specific software stack; it is the responsibility of the providers to investigate, identify and deploy the solution that fits best their individual needs whilst ensuring that the offered services implement the required interfaces and domain languages of the EGI Federated Cloud. The taskforce gives providers a platform to share their implementation solutions for a commonly deployed specific Cloud Management stacks, ensuring a close working relationship with relevant technology providers.

The taskforce activity ensured a close working relationship with relevant technology providers, rapid communication of feedback as well as input on changes needed for deployment, and testing of the new capabilities and services. This effort also ensured that blocking issues needing to be addressed were identified by other areas of EGI (e.g. policy, operations, support and dissemination) and that the testbed, once developed into a production-ready federated e-infrastructure was effectively integrated into the existing EGI infrastructure. The taskforce engaged actively others who were already active in this technology space, enabling diversity in the resource provider community available to research through connection of commercial providers alongside traditionally academic provided resources.

The taskforce activity was organised in six-month phases, with milestones going from the initial setup of the activities to the final integration of the cloud specific service into the EGI e-infrastructure with providers becoming certified members, and thus enabling a production-quality service. As a result of these phases, the taskforce produced a blueprint document for EGI members that wish to securely federate and share their virtualised environments as part of the EGI production infrastructure.

The EGI Federated Cloud was launched in production in May 2014 and can be seen as a new capability offered by the EGI infrastructure, which considerably widens the usage models supported by EGI. Now, web services and interactive applications can be easily integrated in the infrastructure, the computing environments can be finely tuned to satisfy user's needs in term of software (OSs and software packages) and hardware (number of cores, amount of RAM, etc.) and, many solutions are available to store, update and access big amount of data. These new opportunities offered by EGI hugely extended the potential user base of the infrastructure opening the doors to new research communities. The EGI Federated Cloud is already deployed on nearly 20 providers across Europe who together offer more than 6000 CPU cores and 300 TB storage for researchers. This capacity, free at the point of access, is available through IaaS capabilities and interfaces and provide the bases for building higher-level platforms and environments -with the cloud terminology PaaS and SaaS systems– on top. The technologies that enable the EGI cloud federation are developed and maintained by the EGI community, and are based on open standards and open source Cloud Management Frameworks. These technologies are available for institutes and communities who wish to setup their own federated cloud infrastructures.

The EGI Federated Cloud currently supports 26 scientific communities and 50 use cases coming from different scientific disciplines: bioinformatics, physics, earth sciences, basic medicine, arts, language and architecture, mathematics, computer sciences, etc. An extensive programme to support and train users and to facilitate uptake of the new infrastructure within the NGIs is also under development to support the coming community-specific capabilities requested by several research infrastructures from the ESFRI roadmap[5] (BBMRI, EPOS, ELIXIR, DARIAH, EISCAT-3D, INSTRUCT and LifeWatch) planned in the recently started EGI-Engage project[6]. The adoption of the EGI Federated Cloud is ongoing within industry too, through early adopter SMEs from Spain, France and the UK.

2. Related Work

The EGI Cloud Platform seeks to provide heterogeneous research communities with a single set of interfaces to a distributed collection of diverse cloud resources. At the time of the start of the Taskforce there were a number of other projects targeting cloud federations active. Among these, Contrail[7] and Helix Nebula[8] were those comparable to the EGI Cloud effort both in terms of scope and goals. They both targeted multiple European institutions and they both investigated how to provide access to multiple different resource providers by means of a unified interface.

The CONTRAIL project started in 2010 and concluded 2014 and developed a full software stack for building a federation of cloud providers: federated identity management, federated Service Level Agreements (SLAs), a dedicated cloud file system and a Platform as a Service (PaaS) layer. Contrail developments pursued a full-fledge cloud federation distribution following an open source model. However, this method for federation tends to present strong limitations on the sustainability when depending for support and maintenance on a single project with a life span of a few years. By developing a complete new set of software tools instead of 'gluing' existing and self-sustaining components, software development and most of all maintenance overheads become difficult to manage. Limited adoption and lack of maintenance and sustainable development after the end of the project are typical indicators of such issues.

Helix Nebula started in June 2012 and was initially funded through the EC Framework program. It adopts a completely different approach from Contrails: rather than requiring adaptation from the resource providers, it created the federation with the help of a single central broker supporting all of the different proprietary interfaces and cloud instance formats that are part of the federation. This broker, termed the 'Blue Box' within the project, operates as a single interface into all connected resources. Currently, the 'blue box' is a commercial technology based on Sixsq Slipstream[9] that also utilises a unique proprietary interface to provide API access to the broker when not using the web interface. While this approach has the advantage of facilitating the integration of new providers, user communities adopting Helix Nebula are not given the opportunity to easily port their use cases to multiple cloud infrastructures. This is particularly important when considering the diverse landscape of European cloud providers and common scientific requirements of scaling computations to larger, more feature-full infrastructures when needed.

3. The EGI Federated Cloud Architecture

The EGI Federated Cloud architecture is the result of the analysis of the high-level capabilities extracted from users' requirements while taking also into account the needs and expertise of the existing heterogeneous cloud management software locally installed at EGI resource providers. The architecture addresses this challenge by considering each local cloud deployment as an autonomous and abstract subsystem that integrates with the federation through well-defined interfaces (see Figure 1). Wherever possible, open standards are employed for these interfaces to avoid vendor lock-in. Where this was not feasible, community-accepted non-standardised solutions are used. Each resource provider should identify and deploy the solution that fits best their individual needs whilst ensuring that the offered services implement the required interfaces. The taskforce acts as a platform to share implementation solutions and build communities around the different deployed Cloud Management Frameworks.

As shown in the figure, the architecture defines cloud specific capabilities and interfaces and a set of interaction ports with a number of services of the EGI Core Infrastructure[10]. The cloud specific capabilities are: *i*) Virtual Machine (VM) Management and Block storage management, provided with Open Cloud Computing Interface (OCCI)[11,12,13]; *ii*) Data Management, provided with Cloud Data Management Interface (CDMI)[14]; and *iii*) the integration with the Image Management subsystem of EGI, provided with the Open Virtualization Format (OVF)[15] and HEPiX image lists format[16] . EGI Core Infrastructure services integrated are: *i*) Federated AAI, *ii*) Accounting, *iii*) Information System, *iv*) Monitoring, and *v*) Service Registry . The IaaS Cloud capabilities, defined in the architecture, are integrated with the Image Management subsystem, provided as part of the Federated Cloud infrastructure.

The EGI Federated Cloud currently integrates resources from realisations of the abstract cloud management stack for OpenNebula, OpenStack and Synnefo. Whenever possible integration is performed using and contributing to upstream existing projects. New developments are carried out for those specific integration tasks not yet available. The sections below describe each of these capabilities and integration ports with details.

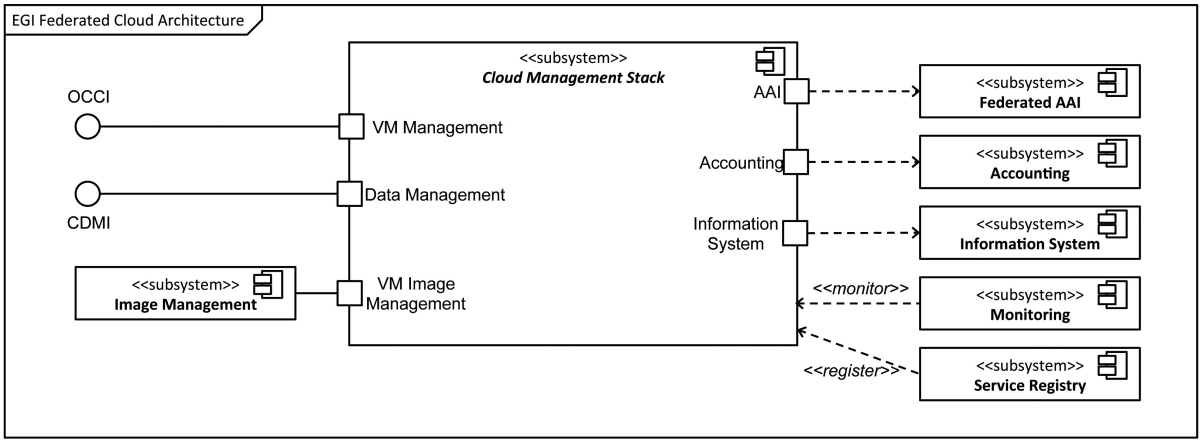

Fig. 1. EGI Federated Cloud Architecture

3.1. Cloud-specific capabilities and interfaces

A cloud provider willing to join the EGI Federated cloud has to implement the following services or interfaces. Providers can either provide VM Management, Data Management or both. If VM Management is provided, integration with the Image Management system is also required.

VM Management. The Open Cloud Computing Interface (OCCI) is a RESTful Protocol and API designed to facilitate interoperable access to, and query of, cloud-based resources across multiple resource providers and heterogeneous environments. The EGI Federated Cloud resource providers implement the Infrastructure specification of OCCI in version 1.1. This has been extended by EGI to support contextualisation and the use of a consistent nomenclature for the resource templates (also known as 'flavours').

Contextualisation is supported by combining the use of cloud-init[17] -a de-facto standard- from within the VM to retrieve contextualisation data, and a FedCloud specific OCCI extension that allows the user to define the data that is to be presented to cloud-init for the contextualisation.

The OCCI implementations are provided as add-on components for the Cloud Management stacks:

- rOCCI[18] for OpenNebula. rOCCI is designed to support multiple back-ends and can be used for providing OCCI in several frameworks, (OpenNebula, CloudStack and AWS).
- OCCI-OS[19] for OpenStack. This already existing OpenStack plug-in was adopted as the OCCI implementation and several bug fixes and extensions are currently contributed by the EGI Federated Cloud members.
- snf-occi[20] for Synnefo. This is a development created for the Synnefo integration into EGI Federated Cloud.

Data Management. Data Management is provided as Object storage and enabled with the support of the Cloud Data Management Interface (CDMI). CDMI defines a RESTful open standard for operations on storage objects. It offers clients a way to operate both on a storage management system (i.e. data containers) and on single data items within a given data container. Each implementation of CDMI can provide a different set of capabilities, which can be discovered by the clients as part of the protocol.

In the EGI Federated Cloud environment, a CDMI reference implementation is provided by Synnefo to support the standard. This implementation can be customised to support different back-ends, the one for Synnefo is provided also. There is ongoing work to develop new back-ends. OpenStack CDMI[21] implementation is also supported in the EGI Federated Cloud.

Image Management. In a distributed, federated Cloud infrastructure users need a way to efficiently manage and distribute VM Images across multiple and heterogeneous Cloud resource providers. The Federated Cloud VM Image

management enables users to register new VM Images and appliances for automatic and endorsed distribution over the federated resource providers. The EGI Applications Database (AppDB) service[22] and HEPiX image lists subscribers at the Resource Providers provide the functionality of this subsystem.

AppDB provides a web interface for users to register and manage VM image metadata and to endorse those VM images as part of Virtual Organisation. The endorsement triggers the subscribers at the providers that, examine and cryptographically verify the provided information, and, pool the new or updated VM Image locally for instantiation. As preferred format for images, the DMTF Open Virtualization Format (OVF) standard is used, which (if needed) subscribers convert to a suitable format supported by the Cloud Management stack.

In order to integrate these pieces, vmcatcher[23] receives and process the list subscriptions and a set of call-out hooks for each specific cloud management stack. The EGI Federated Cloud has developed these hooks for the currently supported stacks.

3.2. Integration with the EGI Core Infrastructure

Apart from the cloud service and interfaces, providers must federate by integrating with EGI core services, which are not cloud-specific but are needed for a successful operation of a production infrastructure.

AAI. Identity federation is one of the core parts of any interoperable infrastructure. Without identity federation, users cannot access the infrastructure transparently, as they would need to posses specific credentials for each of the providers of the infrastructure. EGI controls access to resources by employing personal X.509 certificates and the concept of 'Virtual Organisation' (VO). The VOMS[24] server software acts as an Attribute Authority that issue attributes regarding a VO which are contained in an augmented proxy certificate signed with the user's personal certificate. The proxy and the different attributes determine at the resource provider level if the user is granted to access to the resources. EGI Federated Cloud uses PERUN[25] in OpenNebula and Keystone-VOMS[26] in OpenStack for supporting X.509 proxies and the VOMS extensions included with them.

Accounting. The usage of resources is recorded and integrated in the EGI Core Accounting system. The EGI Federated Cloud has agreed on a Cloud Usage Record (UR) –which inherits from the OGF Usage Record[27]– that defines the data that resource providers must send to EGI's central Accounting repository. A local collector generates records from the Cloud management stacks. UR data is encrypted with the EGI Accounting repository public key and signed with the local configured host certificate's private key, and sent as text payload in secure messages. The EGI Accounting portal provides users, operators and administrators views to access accounting data in a structured and graphical way.

Information Discovery. The EGI information discovery services provide users and tools the means to discover the available resource in the infrastructure. EGI deploys a common information system based on the Berkeley Database Information Index (BDII) with a hierarchical structure distributed over the whole infrastructure. The EGI Federated cloud developed the schema to publish information regarding the available endpoints and their capabilities and an implementation of an information provider[28] that produces the adequate output to be included into the local information providers at each cloud resource provider. AppDB is the main client of the EGI Information service, which is used to display the Resource Providers providing specific VM images and the usage details for instantiating them. This information is also used for operational purposes (the Monitoring subsystem described below depends on these data to be present to operate).

Availability Monitoring. Every service in the EGI infrastructure is monitored via SAM (Service Availability Monitoring). Technology providers must provide specific probes to check functionality and availability of services to ensure compliance with the Operational Level Agreements (OLA) that are put in place once a federated resource provider becomes part of the infrastructure. For the EGI Federated Cloud services there are several monitoring probes available: *i*) Basic "ping" probes to check service reachability *ii*) VM Management probe to check OCCI compliance, and *iii*) Accounting probe to check consistency and freshness of accounting data. CDMI and VM Image Management probes are under development.

Service Registry. EGI's central service catalogue collects the static information of the production infrastructure topology. New service types were defined in the registry to allow providers to expose Cloud resources to the production infrastructure. The catalogue also tracks information for scheduled downtimes, which is important for accurate and correct reliability calculation performed by the Monitoring subsystem.

3.3. PaaS and SaaS services in the EGI Federated Cloud

The EGI Federated Cloud is focused on IaaS services, but its open to external developments that provide higher level services as PaaS and Saas deployed on top of the IaaS capabilities. These services are evaluated by the EGI Federated Cloud taskforce and offered as a way to support new use cases when they fit.

These tools provide users with an alternative way to access the Federated Cloud, hiding the IaaS layer and making simpler the infrastructure usage. Table 1 list the current PaaS and SaaS tools available in the infrastructure. This set is in constant revision and further extended as new developments targeting the EGI Federated Cloud are available.

Table 1. PaaS & SaaS currently available in the EGI Federated Cloud.

Tool	Description
Catania Science Gateway Framework [29]	Visual interface for VMs creation and deployment on different cloud providers based on OCCI standard.
Slipstream	VM image creation and deployment, contextualisation, VM cluster automated deployments, orchestration.
Vcycle [30]	VM lifecycle manager.
COMPSs [31]	VM image creation and deployment, contextualisation, VM cluster automated deployments, orchestration.
VMDIRAC [32]	VM lifecycle manager.
WS-PGRADE [33]	Job, parameter study and workflow submission tool to clouds.

4. Evolution of the EGI Federated Cloud

The EGI Federated Cloud infrastructure continuously evolving in order to satisfy the functionality requirements formulated by the EGI users and providers (e.g. Competence Centers, Virtual Organizations, NGIs, etc.). For the time being, this evolution is being done in three different aspects:

- Exposing additional functionalities from the Cloud Management Frameworks through the standard OCCI interfaces.
- Developing a mechanism for the migration of virtual machines between providers.
- Providing basic brokering capabilities.

The additional required functionalities to be exposed trough the OCCI interface are the support for creating snapshots from a running instance and the support for changing the attached resources to a virtual machine (i.e. resize the instance's CPU, memory, etc.). This functionality is already present in all of the Cloud Management Frameworks being used, however the OCCI standard lacks support for it or there are areas where it is unclear or ambiguous in how a given operation must be performed. The EGI Federated Cloud is working with the Open Grid Forum (OGF) for the inclusion of these extensions or modifications for the upcoming new standard release, that is OCCI 1.2.

Secondly, since the EGI Federated Cloud infrastructure is naturally composed by several resource providers, there is the need to define the set of operations and interactions that a client must fulfil to migrate a suspended virtual machine from one resource provider to another. This design will not only tackle the technical aspects (i.e. transfer

the image, changes in network configuration, etc.) but also will take into account the EGI security and operational policies.

The EGI Applications Database (AppDB) is currently the official catalog of applications and virtual appliances images. It is the central point where users can search for an application or a machine to be executed on the EGI Federated Cloud. Therefore it is logical that it evolves from its current role as a catalog to provide basic VM instantiation capabilities to that users can launch and manage a VM directly in the EGI Federated Cloud from this portal.

5. User support in the EGI Federated Cloud

EGI has developed a support strategy to support the communities interested in exploiting the Federated Cloud capabilities. Providing comprehensive documentation, a low barrier for allowing users to try and evaluate the services, and a well-defined support process to bring use cases towards production were identified as key factors to guarantee a sustainable and successful adoption by future users.

Since the start of production activities in May 2014, EGI has allocated significant effort to create, collect and improve the documentation related to the cloud capabilities, all available on a dedicated Federated Cloud User Support wiki section[34]. The documentation describes the cloud services and includes tutorials and FAQs that guide users to better appreciate the new paradigm and defining step-to-step procedures to understand how to maximise value through using the Cloud. EGI also documents the application models of successful use cases so to facilitate the ease of adoption by future users[35]. Alongside with the documentation, a training program is currently under development[36] covering the usage of the e-infrastructure for new users.

To allow scientists to immediately try and evaluate EGI Federated Cloud in an easy way, we created a catch-all virtual organisation (VO) supported at every provider of the federation. This VO acts as an incubator for new use cases where users can prototype and validate their applications up to six months. The catch-all VO greatly lowered the barrier to entry by removing the need for a large administrative setup process, especially for those users without previous grid experience and therefore lacking the support infrastructure that other communities already have. This is particularly useful understanding whether the infrastructure can be useful for their use cases and also allowed for those communities who have found that the model will work for them, to directly, focus on the integration work without further administrative overheads. In addition, EGI provides a set of Virtual Machine images ready to use and endorsed by EGI Computer Security Incident Response Team (CSIRT) is available in EGI's AppDB to speed-up the testing work without the need of creating any custom images. These VMs are available at all providers and are configured to exploit the main FedCloud capabilities. The set covers commonly used Operating Systems where users' applications can be installed.

The support strategy for any new scientific community interested in the EGI Federated Cloud follows a well-defined workflow of eight phases (see Figure 2 and Table 2) that tries to bring the use cases from the initial tests to a full production status identifying all the activities that should be accomplished during each step.

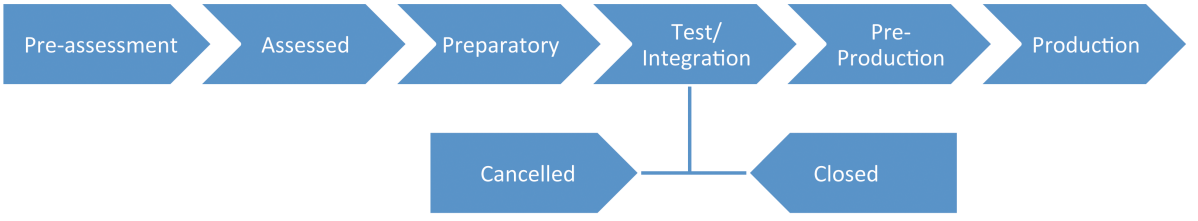

Fig. 2. User Support process workflow

Whenever a new scientific community expresses interest on the EGI FedCloud, the support team (consisting of members of EGI.eu and representatives of technical experts coming from the National Grid Initiatives) organises a face-to-face or web meeting to gather requirements, identify the better strategies to integrate community applications and services and define a complete work-plan. After the first meeting, a team of technical experts is assigned to the communities to continuously track and support them by giving technical aid and evaluating the community progress accordingly to the work plan. Periodic meetings are organised to discuss about the work status and remove hindrances.

Table 2. EGI Federated Cloud use case statuses.

Status	Actions and responsible parties
Pre-assessment	The user support team identified a potential use case that can profit from the EGI Federated Cloud services.
Assessed	The use case requirements are assessed with relevant information added to a dedicated wiki page so to manage the full porting and integration process.
Preparatory	The users and the support team setup the environment to execute tests on the Federated Cloud: configuration of client environment, identification of resource providers, uploading of VM images, etc.
Test and integration	Users access the EGI Federated Cloud through the incubator VO to execute tests while integrating their applications to the infrastructure.
Pre-production	Test and integration are successfully completed. The users create a production virtual organisation or join an existing VO.
Production	The users negotiate and agree on a VO SLA with EGI. EGI agrees on a VO OLA with the resource providers supporting the VO. The use case completed all tests and is regularly making use of the Federated Cloud using a production level VO.
Closed	Test and integration are successfully completed. Use case does not foresee moving into production or the use case ends its life cycle (e.g. all the computations have been completed).
Cancelled	Test and integration are completed unsuccessfully. The user or the support team cancelled the use case.

When the use case successfully completes its integration work and express their interest to start to use the resources in production, the user support team provides assistance: *1*) to obtain a commitment from a set of resource providers to support the community with resources, *2*) to setup a production VO, *3*) to define a Service Level Agreement (SLA) with EGI.eu. The work of the support team is completed when the new community is able to run its services in the EGI production infrastructure in an autonomous way.

6. Conclusions and future work

The development of the EGI Federated Cloud started to support a wider range of user communities and use cases on the e-infrastructure while keeping the investment and expertise already existing at the member organisations of EGI. As a result, the EGI Cloud Infrastructure Platform was defined as a set of extensions and software service add-ons to vanilla deployments of the commonly available Cloud Management Frameworks. These extensions not only provide the interfaces to cloud management and data services but also to the existing and well support EGI Core Platform services, thus providing production-quality federation of e-infrastructure. Open standards for the cloud interfaces assure the necessary interoperability and absence of vendor lock-in that currently exists with many commercial cloud vendors. The selected interfaces –OCCI and CDMI– are implemented on top of all the currently supported frameworks and there is ongoing work to support new platforms. Open formats for the distribution of images (OVF) also contribute to have an interoperable e-infrastructure where VMs can be easily migrated from one provider to another.

The EGI Federated Cloud is currently deployed on nearly 20 providers across Europe who together offer more than 6000 CPU cores and 300 TB storage for researchers and several providers both from Europe and world-wide are in the process of joining the infrastructure. The adoption of the EGI Federated Cloud by research communities has been also significant, not only with those communities which already had established links with EGI, but also with completely new ones. A key part of this is having an easy-to-join prototyping Virtual Organisation where users can get fast feedback on how the services can suit their needs. The well-defined and structured support workflow and the extensive documentation have been shown to be effective with the current existing users, who belong to a wide range of disciplines and come with quite different levels of knowledge of the technology.

Since the initial definition of the architecture of the EGI Federated Cloud, the world of scientific cloud computing has evolved significantly. The EGI community has started to revise the architecture and evolve it according to new user needs in order to remain on the forefront of cloud computing for research and education. The new architecture will allow to create broader collaboration that brings together scientific communities, projects, technology and resource providers to develop and maintain software that enables cloud federations, and to build such federations for diverse user communities. The new EGI Federated Cloud is expected to result in multiple cloud federations, with the currently existing cloud federation becoming only one of them. Core members of the new EGI Federated Cloud collaboration continue to operate and further develop the existing, standard-based EGI cloud federation, but the whole EGI Federated Cloud activity will be open for federated clouds built in different ways using some - or more - of the EGI tools.

Acknowledgements

The authors acknowledge that this work would only have been possible with the effort and contribution from all members of the federated cloud task force. The authors acknowledge the support of the EGI-Engage project [RI-654142] and the use of resources provided by the European Grid Infrastructure.

References

1. ATLAS Collaboration, . Observation of a new particle in the search for the standard model higgs boson with the ATLAS detector at the LHC. *Physics Letters B* 2012;**716**(1):1 – 29. doi:http://dx.doi.org/10.1016/j.physletb.2012.08.020.
2. Sefraoui, O., Aissaoui, M., Eleuldj, M.. Openstack: Toward an open-source solution for cloud computing. *International Journal of Computer Applications* 2012;**55**(3):38–42.
3. Moreno-Vozmediano, R., Montero, R.S., Llorente, I.M.. Iaas cloud architecture: From virtualized datacenters to federated cloud infrastructures. *Computer* 2012;**45**(12):65–72. doi:http://doi.ieeecomputersociety.org/10.1109/MC.2012.76.
4. Synnefo White Paper. Tech. Rep.; synnefo.org; 2014. URL: `https://www.synnefo.org/downloads/synnefo_white_paper_Hi.pdf`.
5. The ESFRI roadmap. 2015. URL: `http://ec.europa.eu/research/infrastructures/index_en.cfm?pg=esfri-roadmap`.
6. The EGI-Engage project - Engaging the Research Community towards an Open Science Commons. 2015. URL: `https://www.egi.eu/about/egi-engage/`.
7. Coppola, M., Dazzi, P., Lazouski, A., Martinelli, F., Mori, P., Jensen, J., et al. The contrail approach to cloud federations. In: *Proceedings of the International Symposium on Grids and Clouds (ISGC'12)*. 2012, p. 1 – 14.
8. Helix nebula project. 2014. URL: `http://helix-nebula.eu`.
9. Sixsq slipstream. 2015. URL: `http://sixsq.com/products/slipstream/`.
10. Drescher, M.. EGI Platform Roadmap. Tech. Rep. MS514; EGI; 2013.
11. Nyren, R., Edmonds, A., Papaspyrou, A., etsch, T.. OCCI core (v1.1). Tech. Rep. GFD-P-R.183; Open Grid Forum; 2011.
12. Metsch, T., Edmonds, A.. Open Cloud Computing Interface - RESTful HTTP Rendering (v1.1). Tech. Rep. GFD-P-R.185; Open Grid Forum; 2011.
13. Metsch, T., Edmonds, A.. Open Cloud Computing Interface - Infrastructure (v1.1). Tech. Rep. GFD-P-R.184; Open Grid Forum; 2011.
14. ISO, . Inormation technology – Cloud Data Management Interface (CDMI). Tech. Rep. ISO/IEC 17826:2012; ISO; 2012.
15. Open Virtualization Format (OVF). 2015. URL: `http://dmtf.org/standards/ovf`.
16. HEPIX Virtualisation Working Group, . Virtual Machine Image List Format. 2015. URL: `http://grid.desy.de/vm/hepix/vwg/doc/html/imagetransfer/imagelist.shtml`.
17. cloud-init. 2015. URL: `https://launchpad.net/cloud-init`.
18. Parak, B., Ustr, Z.. Challenges in achieving iaas cloud interoperability across multiple cloud management frameworks. In: *Utility and Cloud Computing (UCC), 2014 IEEE/ACM 7th International Conference on*. 2014, p. 404–411. doi:10.1109/UCC.2014.51.
19. Occi – OpenStack. 2014. URL: `https://wiki.openstack.org/wiki/Occi`.
20. Asiki, A., Loverdos, C.. snf-occi. 2014. URL: `https://code.grnet.gr/projects/snf-occi`.
21. Cdmi on openstack swift implementation. 2013. URL: `https://github.com/osaddon/cdmi`.
22. EGI Application Database. 2015. URL: `http://appdb.egi.eu`.
23. Synge, O.. vmcatcher. 2015. URL: `https://github.com/hepix-virtualisation/vmcatcher`.
24. Alfieri, R., Cecchini, R., Ciaschini, V., dell'Agnello, L., Frohner, A., Gianoli, A., et al. Voms, an authorization system for virtual organizations. In: *Grid computing*. Springer; 2004, p. 33–40.
25. Denemark, J., Matyska, L., Ruda, M., Jankowski, M., Meyer, N., Wolniewicz, P.. User management for virtual organizations. In: Gorlatch, S., Danelutto, M., editors. *Integrated Research in GRID Computing*. Springer US. ISBN 978-0-387-47656-8; 2007, p. 135–146. doi:10.1007/978-0-387-47658-2_10.
26. Lopez Garcia, A., Fernandez-del Castillo, E., Puel, M.. Identity federation with VOMS in cloud infrastructures. In: *Cloud Computing Technology and Science (CloudCom), 2013 IEEE 5th International Conference on*; vol. 1. 2013, p. 42–48. doi:10.1109/CloudCom.2013.13.
27. Cristofori, A., et al, . Usage record - format recommendation. Tech. Rep. GFD.204; Open Grid Forum; 2013.

28. López, A., Pinto, S., Fernández, E., Parak, B.. cloud-bdii-provider. 2014. URL: https://github.com/EGI-FCTF/cloud-bdii-provider.
29. Ardizzone, V., Barbera, R., Calanducci, A., Fargetta, M., Ingrà, E., Porro, I., et al. The decide science gateway. *Journal of Grid Computing* 2012;**10**(4):689–707. doi:10.1007/s10723-012-9242-3.
30. Vcycle: VM lifecycle management. 2014. URL: http://www.gridpp.ac.uk/vcycle/.
31. Lezzi, D., Rafanell, R., Carrión, A., Espert, I., Hernández, V., Badia, R.. Enabling e-science applications on the cloud with compss. In: Alexander, M., DAmbra, P., Belloum, A., Bosilca, G., Cannataro, M., Danelutto, M., et al., editors. *Euro-Par 2011: Parallel Processing Workshops*; vol. 7155 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg. ISBN 978-3-642-29736-6; 2012, p. 25–34. doi:10.1007/978-3-642-29737-3_4.
32. Méndez, V., Casajús, A., Graciani, R., Fernández, V.. How to run scientific applications with DIRAC in federated hybrid clouds. In: *ADVCOMP 2013, The Seventh International Conference on Advanced Engineering Computing and Applications in Sciences*. 2013, p. 73 – 78.
33. Kacsuk, P., Farkas, Z., Kozlovszky, M., Hermann, G., Balasko, Á., Karóczkai, K., et al. Ws-pgrade/guse generic dci gateway framework for a large variety of user communities. *Journal of Grid Computing* 2012;**10**:601 – 630. doi:10.1007/s10723-012-9240-5.
34. Egi federated cloud user support wiki. 2015. URL: https://wiki.egi.eu/wiki/Federated_Cloud_user_support.
35. Egi federated cloud use cases. 2015. URL: https://wiki.egi.eu/wiki/Federated_Cloud_Communities.
36. Fernández, E., Scardaci, D., Sipos, G., Chen, Y., Wallom, D.. The user support programme and the training infrastructure of the EGI Federated Cloud. In: *Proceedings of 2015 International Conference on High Performance Computing & Simulation (HPCS 2015)*. 2015, p. 9 – 18.