Federating infrastructure as a service cloud computing systems to create a uniform e-infrastructure for research

David C.H. Wallom
University of Oxford, Oxford, UK

Matteo Turilli
Rutgers University, New Jersey, USA

Michel Drescher, Diego Scardaci*
EGI.eu, Amsterdam, NL
* INFN Division of Catania, Catania, Italy

Steven Newhouse
EMBL-EBI, Hinxton, UK

Abstract—**This paper details the state of the art, the design, development and deployment of the EGI Cloud Platform, an e-infrastructure offering scalable and flexible models of utilization to the European research community. While continuing support for the traditional High Throughput Computing model, the EGI Cloud Platform extends its reach to other models of utilization such as long-lived services, and on demand computation. The design and development of the platform was scoped by a set of initial, exemplar user stories describing elemental capabilities of a research cloud infrastructure. Standardized interfaces were integrated and in some cases developed to federate multiple and diverse Infrastructure as a Service (IaaS) cloud resources owned and managed by public and private providers.**

Following a two year period of development, the EGI Cloud Platform was officially launched in May 2014. Since then, the use cases supported have significantly increased both in total number and diversity of model of service required, validating both the choice of enforcing cloud technology agnosticism and of supporting service mobility and portability by means of open standards. These design choices have 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 funding agencies to create economic and social impact from supported research activities.

Keywords—Cloud Computing; e-infrastructure; Federation of Resources.

I. Introduction

Following the successful use of the grid computing platform for the LHC Computing activities [1] and other communities requiring 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. Virtualisation technology and the Infrastructure as a Service (IaaS) cloud service model were considered clear candidates to enable this transformation as many institutes participating in the EGI federation had already invested into virtualisation and IaaS Cloud technology.

An EGI Taskforce was created to evaluate the integration of virtualised resources within the existing EGI production infrastructure. This was to be obtained by provisioning a testbed open to all research user communities, and cataloguing the requirements for community facing services based on or deployed through virtualised resources. 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 and, where open standards were not available, methods that insured broad acceptance in the e-infrastructure community were endorsed.

The Taskforce activity was organised in six-month phases, enabling milestones to be communicated and met before moving onto the next. These milestones alongside their overarching goals are listed below:

1) **Setup** – Identify resource and technology providers and draft the infrastructure model.
2) **Consolidation** – Engage exemplar user communities and start configuration of a testbed.
3) **Federation** – Evolve the testbed into a federated IaaS infrastructure testbed.
4) **Preproduction** – Scope the requirements for both resource providers and core services to reach production.
5) **Integration** – Integrate new cloud specific core services into the EGI e-infrastructure and enforce the processes by which resource providers can become certified members.

The Task Force 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 of a non-technical nature needing to be addressed were identified by other areas of EGI (e.g. policy, operations, support, or 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 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 defining goal of the EGI Cloud Platform is to provide heterogeneous research communities with a single set of interfaces to a distributed collection of diverse cloud resources. This cloud-based research infrastructure can be traced back to the pilot project of the UK National Grid Service [5] developed between 2010 and 2012. Conceived at the time in which open source software to provide public cloud was starting to be publicly available, the UK NGI cloud activities developed along four main dimensions: (i) exploratory comparison of alternative and often competing cloud technologies; (ii) installation, configuration and management of two public cloud infrastructures; (iii) qualitative and quantitative evaluation of the user requirements; and (iv) offering of cloud resources for research and educational use cases.

Three opensource CMF (Rocks Cluster [6], Nimbus [7] and Eucalyptus [8]) were installed on two dedicated systems hosted by the Oxford e-Research Centre and the UK National e-Science Centre. The CMF were then reviewed and compared by focussing on installation requirements, user and administrative interfaces, and overall maturity in terms of reliability and robustness. The comparison outlined the fundamental role played by the cloud interfaces, especially those used to manage virtual machines. These interfaces needed to support both programmatic and visual usage patterns depending on the user community requirements. The reliability of the cloud infrastructure also proved to be of paramount importance. The flexibility offered by the provisioning model resulted in a very fast-paced adoption rate, with the capacity of the two cloud infrastructures reaching saturation after just two weeks of availability.

This project also illustrated the progressing fragmentation of the private cloud distribution landscape. When taking a multi-site offering into account it showed the impossibility of imposing a single distribution on participants even though standardised interfaces were not yet available and implemented. This project also clearly showed how the resource provision model implemented by cloud systems satisfied a diverse set of use cases. Stock and dedicated virtual machines were used by diverse research and teaching communities to satisfy compute and data intensive, alongside service based, use cases. Moving away from the traditional dominance of few types of user communities adopting High Performance Computing (HPC) or grid computing, cloud computing was showing a credible promise to serve the scientific endeavours of the traditionally highly e-infrastructure reliant and also the so called 'long tail of science'.

The findings of this project, and the increase in reliability of software solutions already adopted by members of the EGI community, greatly helped to shape both the vision and operations of the EGI Federated Cloud Taskforce. Alongside the cited technical insights, the UK NGI Cloud project allowed for understanding the requirements imposed by a heterogeneous set of research centres on the concept of resource federation itself. This understanding proved to be as critical as the technical competence required to design a federation architecture serving research activities across the whole Europe.

II. Related Work

At the time of the start of the EGI Federation Cloud Taskforce there were a number of other cloud federation projects active. From among these, Helix Nebula [9] and CONTRAIL [10] were those comparable to the EGI Cloud effort 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.

Initially funded through the EC 7th Framework program (EC FP7), Helix Nebula is an on going activity, started in June 2012. The project adopts a federation method requiring little or no adaptation from the resource providers. Instead, a single central broker is deployed to support all the proprietary interfaces and cloud instance formats offered by the federation. This broker is termed the 'Blue Box' within the project and operates as a single interface into all connected resources. Currently, the 'blue box' [11] is a commercial technology based on SixSq Slipstream [12] exposing a web interface alongside a unique proprietary interface to provide also API access to the broker. While this approach has some merit, especially in terms of homogeneity and integration, by not supporting standard interfaces, APIs and protocols, 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 heterogeneous landscape of European cloud providers and common scientific requirements of scaling computations to larger, more feature-full infrastructures when needed. This amounts to the old foe that cloud was supposed to allow users to move away from vendor lock-in.

Within the CONTRAIL project a different approach was taken, concentrating on interfaces and services necessary to support federation other than the API used to manage virtual appliances within a specific federated cloud solution. Started in 2010 and concluded in 2014, CONTRAIL developed a full software stack offering federation of services managed by multiple cloud providers, federated identity management, federated Service Level Agreements (SLA), a dedicated cloud file system, a Platform-as-a-Service (PaaS) layer, and a unified administrative layer. As such, CONTRAIL can be seen as a full-fledge cloud federation *distribution* developed within the project by following an open source model. Unfortunately, this approach while viable for 'community led' initiatives tends to present strong limitations when depending for support and maintenance on a single project with a life span limited to 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 or sustainable development after the end of the project are typical indicators of such issues.

III. Cloud Federation requirements and scenarios

TABLE 1. CAPABILITIES OF THE EGI FEDERATED CLOUD.

#	Capability	User Stories
1	**Virtual Machine management**	"I want to instantiate a single existing VM image on a remote cloud."
2	**Data management**	"I want to instantiate a VM instance from an image that I have created and is not on the cloud I wish to use." "I want to associate my running VM with a data set in the Cloud." "I want to take snapshots of my running VM for restart purposes."
3	**Integrated information system**	"I want to choose on which resource provider I want to start my single VM." "I need to know about the Virtual Machine Manager (VMM) capabilities the provider offers."
4	**Accounting**	"My usage across different resource providers needs to be recorded and reported to multiple aggregators."
5	**Availability & Reliability**	"Information relating to the availability/reliability and current status of the remote virtualised resource needs to be available to me."
6	**VM & Resource state change notification**	"When the status of the [VM] instance I am running changes (or will change) I want to be told about it."
7	**Integrated AAI**	"I want to use my existing identity, and not re-apply for new credentials to use each component of the service."
8	**VM Image Management**	"I want to use a single VM image across multiple different infrastructure providers."
9	**Brokering**	"I want my VM instance to run on a resource that is suitable based on a set of policies or requirements rather than my choosing directly which resource will run it."
10	**Contextualis ation**	"When I deploy a VM instance on a resource I want to give it configuration information for customisation of the default template. This can only happen when it is up and running."

The model of the federation of Cloud resources are based on high-level capabilities, each illustrated with a set of descriptive requirement statements elicited from user communities, with regard to the utilisation or operation of Cloud related software. This was based on an initial set of six, and subsequently extended to ten, different capabilities each annotated with a user story conveying high level requirements for a federated Cloud infrastructure [13]. Table 1 provides an overview of the capabilities and associated user stories.

VM Management directly implies IaaS Cloud services employing Virtualisation, for which EGI's user communities have strongly voiced their demand. Key aspects here are the implied gain in freedom and control over the software deployed within the Virtual Machine. Though a popular method is to utilise the Virtual Machines storage capability, reference and input data is often stored in publicly accessible repositories, requiring **Data Management** functionality and data repository integration into the VMs. This also covers all other IaaS Cloud functionalities such as VM image storage, snapshots, upload staging area etc.

Integrated information systems provide a mechanism by which EGI user communities and providers, who consume and offer resources from a large number of independent organisations, can have a complete picture of the federated infrastructure services. Furthermore, this capability allows for implementing user-driven load balancing and service availability management for mission-critical projects across federated but otherwise independent resource providers.

Accounting for consumed resources is a fundamental cloud capability, even within a publicly funded IT infrastructure. User communities and infrastructure providers are often engaged in trans-national and international research activities where resources are shared in a 'tit-for-tat' strategy, i.e. an adapted and lightweight peering model for academic infrastructures.

Availability and **Reliability** are vital Key Performance Indicators (KPI) of an infrastructure, helping to build confidence in professional operations and service quality across participating resource providers.

State change notifications for VMs and resources allow for both reactive and proactive management of any set of services deployed on EGI's federated Cloud infrastructure.

Integrated AAI conveys the desire (and, often, the need) of users to use just one, or at most very few authentication and authorisation mechanisms, both across resource providers *and* offered services (not necessarily only Cloud services), often simply because of international and collaborative research activities.

With a firmly established need for a federated (Cloud) infrastructure, it becomes vital for the efficacy of a user's research to ensure that VMs are available for usage at Resource provider when the user needs them. **VM Image management** solves this problem by automating this process as much as is feasible.

Brokering enable users to automate the process of selecting a particular resource provider among a pool of potential and/or accessible resource providers, focusing on higher-level scientific problems. Last but not least, science often requires repetitive and reproducible computations (e.g. simulations, model verifications, sensor data processing, etc.), varying in configuration parameters and input data sets.

Contextualisation is a subset of VM management in that it is the automation of the configuration of a deployed VM instance achieved automatically upon startup or during its lifetime. This may range from trivial activities such as authentication keys being loaded within the VM to highly complex configurations. For example, internal network services or software packages to enable connectivity outside of the deployed instance or the automatic setup of complex deployments made of many VMs.

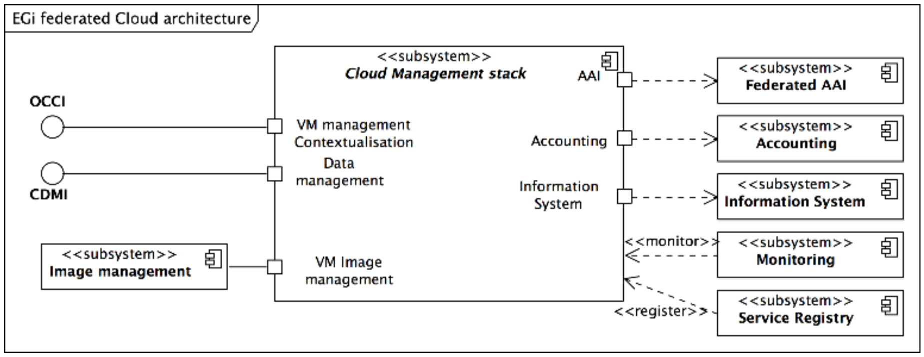

Figure 1. Technical architecture of the EGI Federated Cloud.

A. Technical Architecture

As stated in the introduction, the federation model and architecture need to take into account not only the user requirements illustrated in an earlier section, but also the needs, assets and expertise already acquired in the existing EGI resource provider community in a heterogeneous, locally maintained Cloud software landscape.

The technical architecture addresses this by modelling local Cloud deployments as *abstract subsystems* integrating with the federated Cloud infrastructure through well-defined interfaces (Figure 1). Wherever possible and feasible, publicly defined specifications such as Open Cloud Computing Interface (OCCI) [14, 15, 16], Cloud Data Management Interface (CDMI) [17], Usage Records (UR) [18] or Grid Laboratory Uniform Environment (GLUE) [19] are employed. Where this was not feasible, popular and community-accepted non-standardised solutions were adopted.

The capabilities are mapped onto a number of subsystems of the EGI Cloud Platform (Figure 1), of which some are shared with other ICT platforms deployed in the EGI production infrastructure: Federated AAI, Accounting, Monitoring, Information System, and Service Registry (the EGI Platform Roadmap [20] describes this in more detail). The Cloud Management stack subsystem is modelled to implement VM Management, Contextualisation and Data Management, and to integrate with the Image Management subsystem provided as part of the Federated Cloud infrastructure. Where required, ports model the need for a Cloud Management stack realisation to integrate with surrounding subsystems, components, and specifications.

EGI encourages and facilitates resource providers deploying the same underpinning Cloud Management frameworks (see Table 2) to form lightweight interest groups within EGI providing help to each other in the integration activities, share integration code development and maintenance effort, and other overlapping interests.

B. Cloud-specific capabilities and interfaces

Federating cloud services requires service providers conforming to common and, where possible, open standard interfaces describing several key functions and capabilities. Resource providers have to fully implement all of the following services or interfaces to be fully compliant with the EGI Cloud Platform.

1) VM Management & Contextualisation: OCCI

The IaaS management interface, as specified by OCCI, models Cloud resources and capabilities in a flexible and extensible way. OCCI consistently defines a RESTful rendering for all its models and externally defined extension. OCCI also defines a model for IaaS services in version 1.1.

The presented model defines an extension to OCCI supporting VM templates in terms of combining a set of compute resource attribute settings (e.g. number of cores) using a descriptive term. Users of the federated Cloud services can use these templates (also called “flavours”) consistently across all federated resource providers.

Contextualisation is supported in a standardised manner by combining the use of cloud-init [21] from within the VM to retrieve contextualisation data, and a specific OCCI extension that allows the user to define the data that is to be presented to cloud-init for the contextualisation.

TABLE 2. OVERVIEW OF AVAILABLE INTEGRATION FOR DEPLOYED CLOUD MANAGEMENT FRAMEWORKS.

	Integration					
Cloud Management Stack	Fed. AAI	Monitoring	Accounting	Img. Mgmt.	OCCI	CDMI
OpenStack	Y	Y	Y	Y	Y	Y
OpenNebula	Y	Y	Y	Y	Y	Y
Synnefo	Y	Y		Y	Y	Y
WNoDeS	Y				Y	
StratusLab	Y	Y	Y		Y	
Cloudstack					Y	
Emotive	Y					Y

None of the currently supported CMFs utilise the OCCI interface as their de facto native service but expose the OCCI services through available add-on components as intermediate layers. For Openstack this is OCCI-OS[22], snf-occi for Synnefo [23] and rOCCI[24] for OpenNebula, StratusLab and Cloudstack. rOCCI is also used for a number of other CMFs though within the federated cloud these are currently not used. These realisations implement and expose VM management and

VM Contextualisation capabilities through a single port, allowing users to manipulate and manage VM instances (almost) without restriction.

2) Data management: CDMI

Data management is enabled by supporting CDMI, which defines a RESTful open standard for operations on storage objects. CDMI 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. Since CDMI allows for partial implementations, the exact level of support depends on the concrete implementation and is exposed to the client as part of the protocol.

CDMI is built around the concept of objects, which vary in supported operations and metadata schema. Each Object has an ID, which is unique across all CDMI deployments. There are four objects most relevant in the context of the EGI cloud federation: (i) 'data object', an abstraction for a file with rich metadata; (ii) 'container', an abstraction for a folder; (iii) 'capability', to expose information about the feature set of an object; and (iv) 'domain', to expose deployment specific information.

Currently, the reference implementation of CDMI provided by Synnefo is utilized within the EGI federated cloud but a further implementation from the Barcelona Supercomputing Centre, named EMOTIVE [25], is also available.

3) VM Image management: S/MIME, JSON, X.509, SHA512

CMFs must integrate with the EGI VM Image management subsystem enabling users to register new VM Images and appliances for automatic distribution over the federated resource providers. The system provides a web interface for the user to notify supporting resource providers of the existence of a new or updated VM Image. These then examine the provided information, and, pending their decision, pool the new or updated VM Image locally for instantiation. This system offers the following features:

- VM Image lifecycle management – Scalable Software Lifecycle Management across EGI.
- Automated VM Image distribution – Publish VM image changes triggering resource providers to obtain the updates.
- Asynchronous distribution mechanism – Scalable solution to distribute VM image updates efficiently.

The concrete realisation of the VM catalogue is obtained through the Application Database [26] service within EGI, although additional catalogues could be connected to the infrastructure to serve, for example, specific communities. The VM image management system provides also a notification mechanism to resource providers (named "vmcaster"), and a service to receive and process these notifications (named "vmcatcher")[27]. Vmcatcher is deployed as part of the integration process for a local Cloud management stack by configuring the vmcatcher call-out hooks.

Image lists are expressed as JSON [28] documents for easy construction and parsing. As a text-based information rendering, it fits very well as payload with S/MIME [29] messages that are cryptographically signed using SHA512 [30] as message digest algorithm, and X.509 [31] certificates as authentication tokens for signer, sender and receiver of image lists.

C. Integration with EGI's core infrastructure

As already indicated, the federated Cloud infrastructure makes use of a number of existing core EGI services. While being essential for the successful operation of a federated production infrastructure, they are not Cloud-specific and are of less interest in the context of this paper thus they will be only concisely described.

1) AAI and Virtual Organisation management: PKI, proxy certificates

Access to IT resources in EGI is controlled through employing a federated authentication infrastructure operated by national e-Science Certificate Authorities (CAs) outside of EGI's direct operational control as a Public Key Infrastructure (PKI) [32]. Upon appropriate offline identification, users will be issued with a personal X.509 certificate. Users must be member of a group of researchers (academic, commercial or mixed) called "Virtual Organisation" (VO) which acts as an Attribute Authority issuing VO attributes that must be contained in an augmented proxy certificate signed with the user's personal certificate. Only when presenting this proxy certificate, users may be granted access to the resources as gatekeeping services check the certificate subject's DN as well as the VO attributes.

The cloud federation uses a tool named PERUN [33] to manage virtual organisations. . PERUN offers quick and simple management services for each VO including interfaces that simplify the complex process of user management. The support for the management of virtual organisations within the CMF has been achieved by means of both a native capability implemented within rOCCI and the integration of a federated identity capability in the OpenStack Keystone service [34].

2) Accounting: Messaging, STOMP, Usage Records (UR)

Cloud Management stack realisations must record usage of exposed resources into a Usage Record format. Integration activities require the service providers to deploy dedicated code that harvests and sifts through local data sources to generate UR data. Implementations exist for OpenNebula, OpenStack and Synnefo. Certified federated resource providers must install an accounting probe provided by the EGI accounting system that collects the generated UR data, and sends them to a central accounting database using the EGI messaging infrastructure. UR data is encrypted with the Accounting DB's public key, signed with the locally configured host certificate's private key, and then sent as text payload in secure STOMP [34] messages. A portal allows users, operators and administrators to access accounting data in a structured way.

3) Information System: Berkeley DB, LDAP, LDIF, GLUE

The EGI Information System is a distributed system of hierarchical information indexes based on the Berkeley DB [36] technology, and is commonly named BDII. Deployed in a three-layer fashion, top-BDIIs are deployed at a national level while aggregate data are collected from subordinate local BDIIs installed at the resource provider level. Finally, a local BDII collect data from each Cloud Management stack. This third layer is called the "information provider" and is the integration target every Cloud Management Framework should provide. Information is modelled and stored in GLUE format throughout the system, from information provider to every top-BDII. All BDIIs provide a query interface using LDAP v3 [37] and GLUE's LDAP rendering. Data updates sent between BDII instances are encoded using LDIF [38].

The Information system is also used for operational purposes. For example, the Monitoring subsystem depends on the Information System for certain data to be present in order to properly operate.

4) Monitoring: Messaging, Nagios

EGI's monitoring system is used to monitor both operational services as well as shared resources. Consequently, the services of the federated Cloud infrastructure are monitored in a variety of ways ensuring both user satisfaction and compliance with the Operational Level Agreements (OLA) that are put in place once a federated resource provider becomes operational and certified. A number of central and regional/national customised Nagios [39] instances are operated, and predefined sets of Nagios-based tests are deployed centrally and nationally. A service availability calculation engine receives test results via EGI's messaging infrastructure and regularly produces predefined reports. Multiple tests are in place for cloud services: (i) the VM Management probe checks OCCI compliance; (ii) the VM Management "ping" probe checks basic service reachability; (iii) the Data Management "ping" probe checks basic service availability; (iv) the Accounting probe checks consistency and freshness of accounting data; and (v) the Information System probe checks resource providers publish up-to-date technical data. More tests are continuously developed and added to the roadmap where need arises.

5) Service registry

EGI's central service catalogue is used to catalogue both the static information of the production infrastructure topology and their current state. The service registry collects information about services in testing mode, production, whether they are monitored, and whether they are certified. It also tracks information for scheduled downtimes, which is important for accurate and correct reliability calculation performed by the monitoring subsystem.

IV. User Communities and Supported Usage Models

Since production launch in May 2014, EGI has allocated significant effort to developing and refining a suitable strategy that supports the many communities interested to exploit this new capability. EGI has been also documenting the application models of successful use cases so to facilitate the ease of adoption by future users. The need for clear and complete documentation, an easy way to allow scientists to immediately try and evaluate the Cloud, and a well-defined support process to bring use cases towards a production status were identified as key points to guarantee the success of the EGI Federated Cloud. Thus, the federation User Support wiki [40] was created, collecting and improving the documentation related to the new Cloud services, creating tutorials that guide users to better appreciate the Cloud paradigm and defining step-to-step procedures to understand how to maximise value through using the Cloud.

We decided to create a catchall virtual organisation (VO), acting as an incubator for any new use cases, reducing the time between the first contact with a community and their first test of the cloud e-Infrastructure. This catchall VO has been enabled in all certified sites, to engage all resource providers in this effort. The stated mission of the VO is to support application prototyping and validation for up to six months. The catchall VO greatly lowered the barrier to entry by removing the need for a large administrative setup process. This is particularly useful for those prospective user communities that mainly want to understand whether the infrastructure can be useful for their use cases. This 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.

The user support process includes a few manual labour intensive steps, therefore it was essential to define a clear process to bring the use cases from the initial tests to a full production status. This end-to-end process consists of eight phases defining actions and responsible parties for each phase as shown in Fig. 2.

- **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.).

Figure 2. Support workflow for the full use-case life cycle.

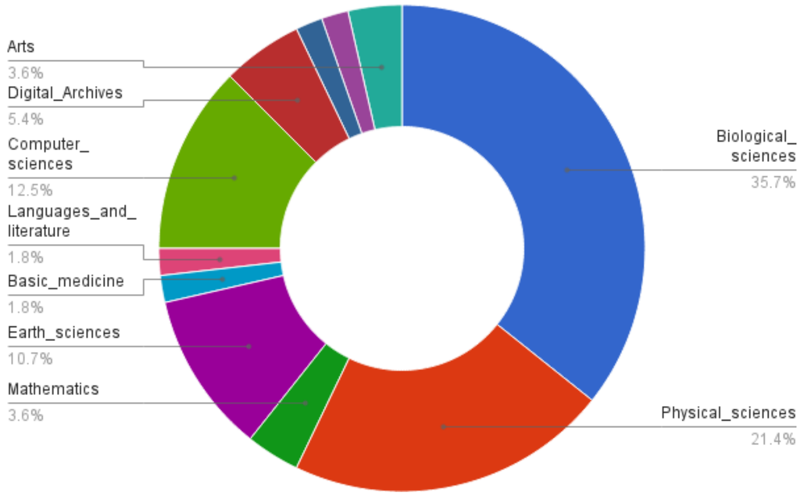

Figure 3. Distribution of use cases amongst domains.

- **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 organization or join an existing VO.
- **Production**: the users negotiate and agree on a VO SLA with EGI. EGI agree 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.

A. Supported communities

As shown in fig. 3 we are currently supporting 26 communities and 50 use cases coming from different scientific disciplines (Figure 3) [41]. Of these, 11 are in full production (including the integration of some high level tools in the Federated Cloud) and 8 are in pre-production.

A third of the Federated Cloud use cases come from the biological sciences. Furthermore, the Swedish (BILS) [42] and Finnish (CSC) [43] ELIXIR nodes are working at integrating some of their services on the Federated Cloud and the BioVel [44] project is running four applications in production. A further example of biological applications is the READemption [45] use case, a next generation sequencing application developed by the Würzburg University in Germany.

Physical sciences represent about a fifth of the use cases with possibility of expansion as CERN's ATLAS [46], CMS [47] and LHCb [48] communities evaluate running their jobs on the Federated Cloud.

Earth sciences use cases include large communities as DRIHM [49] and VERCE [50] running, respectively, hydrology and hydraulic models, and seismology applications. The well-known WRF application has been integrated by the CHAIN-REDS [51] project and Jena University successfully tested the Federated Cloud to run JAMS [52], a framework to build up complex hydrological models.

In addition, some private companies showed interest to the EGI Federated Cloud services. For example, Engineering, an Italian SME, is currently running in the EGI Federated Cloud the HAPPI toolkit, a platform for data preservation developed for the SCIDEP-ES [53] project, and the backend services of INERTIA [54], an application to retrieve energetic data of final-occupant of tertiary building.

The user base of the Federated Cloud is not yet complete. We are currently supporting new and developing use cases from basic medicine, arts, language and architecture, mathematics and computer sciences. It is interesting to note that the biological sciences are emerging as the most dominant in terms of numbers and early adoption of cloud resources, a departure from the dominance of the EGI High Throughput Platform by the Physics community.

B. Use case models

The EGI Federated Cloud is able to support different use case models by design. In this section, we briefly describe some of the most relevant, highlighting how each benefit from Federated Cloud services in different yet typical ways.

1) READemption - Heavy computation and large memory

READemption is a pipeline for the computational evaluation of RNA-Sequence data. It was originally developed to process dRNA-Seq reads [55] originating from bacterial samples. It has since been extended to process data generated in different experimental setups and from all domains of life. The functions cover read processing and aligning, coverage calculation, gene expression quantification, differential gene expression analysis as well as visualization.

The READemption pipeline is a typical example of computationally intensive and large memory applications. It requires VMs with up to 20 cores each and up to 70 GB of RAM. Furthermore, 1 TB of storage is attached to each VM by means of the block storage capability.

The READemption community already has a computational infrastructure but external cloud resources are required when data bursts need to be processed. In the past, Amazon Web Services [56] were used to manage these peaks of demand but Amazon billing model proved to be too expensive. For this reason, the READemption community first tried the EGI Federated Cloud and eventually decided to use it to manage the community's peak demand.

2) BioVel - Web Services

BioVeL is a virtual e-laboratory that supports research on biodiversity issues using large amounts of data collected from cross-disciplinary sources. BioVeL offers the possibility to use computational "workflows" to process data. BioVeL is an EC FP7 funded project.

The BioVel community has adopted the EGI Federated Cloud to deploy web services implementing the workflows and the BioVel Portal, a browser-based interface that allows for the BioVel users to upload, run and manage workflows, as well as the viewing and downloading of their outputs. The BioVel web-services currently available in the Federated Cloud are OpenRefine, OpenModeller, and BioStif. This set will be expanded in the next months.

3) Peachnote - Heavy computation

Peachnote [57] is a music score search engine and analysis platform and is one of the private sector use-cases operating within the Federated Cloud. The system is the first of its kind and can be thought as the analogue of Google Books but for music scores. Hundreds of thousands of music scores are being digitized by libraries all over the world but in contrast to books, they generally remain inaccessible for content-based retrieval and algorithmic analysis. Thanks to its search engine and Ngram viewer, Peachnote provides visitors and researchers access to a large and growing amount of symbolic music data.

Building up the corpora of musical sheets requires significant amount of computing capacity. Millions of sheets need to be converted from JPG images to music XML, sometimes multiple times, using different configurations of the converter algorithm. Peachnote is integrated with the Simple Queue Service (SQS) of Amazon. Amazon SQS is the coordinator of computation, with sites of the EGI Federated Cloud used for the back-end calculations.

4) Engineering - SME

Engineering [58] is an Italian SME involved in research projects in many different sectors: engineering, cultural heritage, agriculture, computer science, smart cities etc. Initially they tested the Federated Cloud by deploying two services: the SCIDIP-ES HAPPI toolkit and the backed system of the INERTIA platform. The HAPPI toolkit, developed for the FP7 SCIDIP-ES project, is a tool for cultural heritage and data preservation that supports an archive manager and curator to capture and manage part of the Preservation Descriptive Information (PDI). INERTIA is a project addressing the "structural inertia" of existing Distribution Grids. It adapts the Internet of Things/Services principles to the Distribution Grid Control Operations thanks to the introduction of new active elements combined with the necessary control and distributed coordination mechanisms. . The INERTIA platform is made of a back-end service and an iOS application and it allows for retrieving energetic data of final-occupant of tertiary building.

These two test cases were successfully completed and the HAPPI toolkit and the INERTIA back-end services deployed on the EGI Federated Cloud are now used as main instances to support Engineering's users. Furthermore, Engineering has expressed interest to investigate the deployment of further services in a near future.

5) CHAIN-REDS - Infrastructure & Application broker

CHAIN-REDS is an EC FP7 project aiming at promoting and supporting technological and scientific collaboration across geographically distributed e-Infrastructures. The goal is to help define a path towards a global e-Infrastructure ecosystem.

CHAIN-REDS offers to its users an easy-to-use Science Gateway [59] encapsulating many scientific applications and allowing them to deploy virtual services on supported resources. The CHAIN-REDS community decided to integrate its science gateway with the EGI Federated Cloud to execute scientific applications by means of VMs dynamically instantiated on the cloud. In this way the virtual services started by the users are deployed on the VMs of the EGI Federated Cloud. The CHAIN-REDS science gateway acts as an application broker (SaaS) for the Federated Cloud supporting application instance deployment across multiple connected resources to which the user is given access.

V. FURTHER WORK

Following the successful launch of the production federated cloud services there are a number of key additional activities and capabilities that need to be supported to respond to the feedback of user communities. The first of these are at an infrastructure level with the introduction of capabilities that enable the multi provider federation model to be exploited more completely as well as to cope with situations or scenarios typical of a production system, that hitherto were not encountered during initial design and development. As with contextualisation, a number of capabilities need further investigation; indeed, they are considered standards within current public cloud architectures but are not enabled by default in a uniform manner across our cloud management frameworks or within their standards based interfaces. Key amongst these is the need to support complex network configurations inside and between individual cloud resources providers. Examples are the allocation of IP addresses, network groups and VPN services as well as the setup of dedicated network pipes between resources. Some of these will require further extensions to the OCCI standard, but some others will depend on, for example, whether and how Software Defined Networking (SDN) implementations will be integrated. In this context, a particularly challenging element will be establishing boundaries between an individual cloud system and the connecting network.

Secondly, there is the need to look at the business and operating models under which the federated cloud can operate going forward. It is essential that EGI continues broadening both the number and type of resource providers and user communities from all sectors, public and private. Alongside the strengthening of the underpinning platform, we aim to further develop the operating and business models the federated cloud is able to support. Building upon the work of the EGI task force on pay per use models [60], updates to core infrastructure services to support pricing structures has already been started. This includes requirements on methods and tools to support the automated matching of providers and consumers within the federation. This is essential as in a large number of cases the implemented model of federation predicates a weaker connection between the consumer and provider than in previous e-infrastructure services (e.g. grid),.

To support future use cases where there may be contractual obligations on the provision of resources and their operational quality, the EGI has been working on integrating the FitSM service management framework [61] into its offering though this work is still in its an early stages when considering the provision of cloud resources..

Finally, to further support user communities it will be important to broaden the number of different high-level tools available. We have seen, as described in section IV, that different existing communities have explored a wide range of application design models and utilisation schemes through the access tools currently provided. Thanks to the integration of other popular management and provisioning tools as front ends, we will deliver an offering more easily integrated with other services from within and without the EGI ecosystem. It is important though that this is achieved through the support for standards rather than the design of yet another management interface as clearly demonstrated with the deployment of the current infrastructure. Important is also avoiding the maintenance of a large number of shims between services and the provisioning systems as they would generate significant overheads for both the service providers and the users, should they wish to migrate to another platform.

VI. Conclusion

The Federated Clouds Task Force started exploring a federation of private institutional Cloud deployments to support a wider variety of user communities in their need for e-infrastructure resources. The developed EGI Cloud Platform describes the extensions and software service add-ons for standard installations of different IaaS Cloud Management Frameworks (CMF) that are already deployed by EGI members with varying degree of popularity. This approach to cloud federation allows for the separation of concerns on which CMF to deploy, from the mechanics of joining the EGI Federated Cloud. We have supported not only the definition of the technical interfaces necessary to connect to cloud management and data services but also existing and well supported EGI Core Platform services. This allows leveraging years of experience in the federation of e-infrastructure services at all levels accumulated by EGI. Following the success to date of the EGI Federated Cloud in attracting resource providers, the e-infrastructure expansion program is continuing. Therefore, we are currently investigating the integration of a higher number of heterogeneous platforms whilst maintaining the functionality and capabilities as with currently deployed technologies. A number of on going pilot projects with research communities stemming from both within and outside the current EGI ecosystem have demonstrated that a cloud-based e-infrastructure is better suited to support a greater diversity of research requirements than other types of e-infrastructure currently available. This is allowing a significant growth in the number of effective users that are being supported, leading to the integration of different models of utilisation within the cloud-based e-infrastructure. Thanks to this integration process, we are building a catalogue of operational and application design models with which we can engage further communities and respond to their needs.

We assessed the current set of recommended cloud standards by deploying their implementations within the e-infrastructure and evaluating their design methodologies. In many cases, the assessment process led towards further work with the standardization working groups to ensure that the next versions of these standards will natively support our requirements. As such, we consider this as one of the greatest successes of the federated cloud, showing that cloud standards are both useful and essential when multiple providers need to be able to support a large number of use cases while staying competitive within the cloud market place.

Acknowledgments

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; Alison Packer, Álvaro López García, Alvaro Simon, Binh Minh Nguyen, Björn Hagemeier, Boris Parak, Boro Jakimovski, Cal Loomis, Daniele Cesini, Daniele Lezzi, David Blundell, Elisabetta Ronchieri, Emir Imamagic, Enol Fernandez, Feyza Eryol, Florian Feldhaus, Gergely Sipos, Jan Meizner, John Gordon, Kostas Koumantaros, Malgorzata Krakowian, Marios Chatziangelou, Marco Verlato, Marica Antonacci, Mattieu Puel, Michel Jouvin, Miroslav Ruda, Nuno L. Ferreira, Owen Synge, Paul Miller, Peter Solagna, Piotr Kasprzak , Roberto Barbera, Ruben Valles, Sándor Ács, Salvatore Pinto, Silvio Spardi, Soonwook Hwang, Stuart Pullinger, Thijs Metsch, Tomasz Szepieniec, Viet Tran, Zeeshan Ali Shah. The authors acknowledge the support of the EGI-Inspire project [RI-261323] 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, Volume 716, Issue 1, 17 September 2012, Pages 1-29, ISSN 0370-2693, http://dx.doi.org/10.1016/j.physletb.2012.08.020.
2. Sefraoui, O., Aissaoui, M., & Eleuldj, M. (2012). OpenStack: toward an open-source solution for cloud computing. *International Journal of Computer Applications*, *55*(3), 38-42.

3. Sotomayor, B., Montero, R. S., Llorente, I. M., & Foster, I. (2009). Virtual infrastructure management in private and hybrid clouds. *Internet computing, IEEE*, *13*(5), 14-22.
4. Synnefo Open-source IaaS Whitepaper, http://www.synnefo.org/downloads/synnefo-white-paper-Hi.pdf
5. Geddes, N. (2006, December). The National Grid Service of the UK. In *e-Science* (p. 94).
6. Sacerdoti, F. D., Chandra, S., & Bhatia, K. (2004, September). Grid systems deployment & management using Rocks. In *Cluster Computing, 2004 IEEE International Conference on* (pp. 337-345). IEEE.
7. Keahey, K., Figueiredo, R., Fortes, J., Freeman, T., & Tsugawa, M. (2008). Science clouds: Early experiences in cloud computing for scientific applications. *Cloud computing and applications*, *2008*, 825-830.
8. Nurmi, D., Wolski, R., Grzegorczyk, C., Obertelli, G., Soman, S., Youseff, L., & Zagorodnov, D. (2009, May). The eucalyptus open-source cloud-computing system. In *Cluster Computing and the Grid, 2009. CCGRID'09. 9th IEEE/ACM International Symposium on* (pp. 124-131). IEEE.
9. Helix Nebula project http://helix-nebula.eu
10. Coppola, Massimo, et al. "The Contrail approach to cloud federations." *Proceedings of the International Symposium on Grids and Clouds (ISGC'12)*. Vol. 2. 2012.
11. Helix Nebula Marketplace Catalogue:. http://www.helix-nebula.eu/sites/default/files/Helix%20Nebula%20Marketplace%20Catalogue_0.pdf
12. SixSq SlipStream, [Online] Available: http://sixsq.com/products/slipstream.html
13. D Wallom, M Drescher, "MS521: EGI Federated Cloud Blueprint 2nd edition", June 2014 https://documents.egi.eu/document/2091
14. R. Nyren, A. Edmonds, A. Papaspyrou, and T. Metsch, "Open Cloud Computing Interface - Core," GFD-P-R.183, April 2011. [Online]. Available: http://ogf.org/documents/GFD.183.pdf
15. T. Metsch and A. Edmonds, "Open Cloud Computing Interface - RESTful HTTP rendering," GFD-P-R.185, June 2011. [Online]. Available: https://www.ogf.org/documents/GFD.185.pdf
16. T. Metsch and A. Edmonds, "Open Cloud Computing Interface - Infrastructure," GFD-P-R.184, April 2011. [Online]. Available: http://ogf.org/documents/GFD.184.pdf
17. SNIA Technical Position, "Cloud Data Management Interface (CDMI) 1.0.2", June 2012 http://snia.org/sites/default/files/CDMI%20v1.0.2.pdf
18. A. Cristofori, et al, "Usage Record – Format Recommendation", July 2013, http://www.ogf.org/documents/GFD.204.pdf
19. S. Andreozzi et al, "GLUE Specification v. 2.0", March 2009, http://www.ogf.org/documents/GFD.147.pdf
20. M Drescher, "MS518 EGI Platform Roadmap", June 2014, https://documents.egi.eu/document/2232
21. S Moser et al, "Configuration and customization of cloud instances", https://cloudinit.readthedocs.org/en/latest/index.html
22. T. Metsch, A. Edmonds. (2012). OCCI-OS [Online] Available: https://wiki.openstack.org/wiki/Occi [Accessed: September 29, 2014].
23. A. Asiki, Ch. Loverdos. (2014). snf-occi [Online] Available: https://code.grnet.gr/projects/snf-occi [Accessed: September 29, 2014]
24. Parák, B., Šustr, Z., Feldhaus, F., Kasprzakc, P., & Srbac, M. (2014, March). The rOCCI Project–Providing Cloud Interoperability with OCCI 1.1. In *International Symposium on Grids and Clouds (ISGC)* (Vol. 23, No. 28).
25. Goiri, I., Guitart, J., Torres, J.: Elastic Management of Tasks in Virtualized Envi-ronments. In: XX Jornadas de Paralelismo, JP 2009, A Corua, Spain, September 16-18, pp. 671–676 (2009)
26. EGI Applications Database, "AppDB," https://appdb.egi.eu/.
27. O. Synge, et al., "Virtual machine image transfer," in HEPIX Virtualisation Working Group, 2012. [Online]. Available: http://grid.desy.de/vm/ hepix/vwg/doc/pdf/Book-a4.pdf
28. JSON (JavaScript Object Notation), the a lightweight data-interchange format, http://json.org/
29. Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.2, Message Specification, January 2010, http://tools.ietf.org/html/rfc5751
30. Descriptions of SHA-256, SHA-384, and SHA-512, http://csrc.nist.gov/groups/STM/cavp/documents/shs/sha256-384-512.pdf
31. Internet X.509 Public Key Infrastructure, Certificate and CRL Profile, January 1999, https://www.ietf.org/rfc/rfc2459.txt
32. Vacca, Jhn R. (2004). *Public key infrastructure: building trusted applications and Web services*. CRC Press. ISBN 978-0-8493-0822-2.
33. Denemark, J., Matyska, L., Ruda, M., Jankowski, M., Meyer, N., & Wolniewicz, P. (2007). User Management for Virtual Organizations. In *Integrated Research in GRID Computing* (pp. 135-146). Springer US.
34. Lopez Garcia, A.; Fernandez-del-Castillo, E.; Puel, M., "Identity Federation with VOMS in Cloud Infrastructures," *Cloud Computing Technology and Science (CloudCom), 2013 IEEE 5th International Conference on* , vol.1, no., pp.42,48, 2-5 Dec. 2013
35. Simple Text Oriented Messaging Protocol 1.2, October 2012, http://stomp.github.io/
36. The Oracle Berkeley DB, https://launchpad.net/berkeley-db
37. Lightweight Directory Access Protocol (LDAP): Technical Specification Road Map, RFC4510, https://tools.ietf.org/html/rfc4510
38. The LDAP Data Interchange Format (LDIF) - Technical Specification, RFC 2849, http://tools.ietf.org/html/rfc2849
39. Nagios: IT Infrastructure monitoring system, http://www.nagios.org/
40. EGI Federated Cloud User support: https://wiki.egi.eu/wiki/Federated_Cloud_user_support
41. EGI Federated Cloud use cases: http://go.egi.eu/fcc
42. Bioinformatics Services for Swedish Life Science (BILS): https://bils.se/
43. CSC - Finnish IT Center for science: https://www.csc.fi/
44. BioVel - Biodiversity Virtual e-Laboratory: http://www.biovel.eu/
45. *READemption – A tool for the computational analysis of deep-sequencing-based transcriptome data*. Konrad U. Förstner, Jörg Vogel, Cynthia M. Sharma. 2014, Bioinformatics, Aug 13. Fulltext, Pre-print at bioRxiv.
46. CERN - ATLAS experiment: http://atlas.ch/
47. CERN - CMS experiment: http://cms.web.cern.ch/
48. CERN - The LHCb collaboration: http://lhcb.web.cern.ch/lhcb/
49. Distributed Research Infrastructure for Hydro-Meteorology Study (DRIHMS) project: http://www.drihms.eu/
50. Virtual Earthquake and seismology Research Community in Europe e-science environment: http://www.verce.eu/index.php
51. Coordination and Harmonisation of Advanced e-Infrastructures for Research and Education Data Sharing (CHAIN-REDS): http://www.chain-project.eu/
52. JAMS - Jena Adaptable Modelling System: http://jams.uni-jena.de/
53. SCIDIP-ES HAPPI toolkit: http://int-platform.digitalpreserve.info/dashboard/happi-toolkit/
54. The INERTIA project: http://www.inertia-project.eu/inertia/
55. Sharma CM, Hoffmann S, Darfeuille F, Reignier J, Findeiss S, Sittka A, Chabas S, Reiche K, Hackermüller J, Reinhardt R, Stadler PF, Vogel J. The primary transcriptome of the major human pathogen Helicobacter pylori. Nature. 2010 Mar 11;464(7286):250-5. doi: 10.1038/nature08756. Epub 2010 Feb 17.
56. Amazon Web Services Cloud Computing Services: http://aws.amazon.com/
57. Peachnote: http://www.peachnote.com
58. Engineering web site: http://www.eng.it/home/home-eng.dot
59. Valeria Ardizzone, Roberto Barbera, Antonio Calanducci, Marco Fargetta, E. Ingrà, Ivan Porro, Giuseppe La Rocca, Salvatore Monforte, R. Ricceri, R. Rotondo, Diego Scardaci, Andrea Schenone: "The DECIDE Science Gateway". Journal of Grid Computing Vol. 10(4), pages 689-707, Editor Springer 2012.
60. EGI Pay for Use Proof of Concept: https://wiki.egi.eu/wiki/EGI_Pay-for-Use_PoC:Home
61. FitSM, "Standard for lightweight service management in federated IT infrastructures", http://www.fedsm.eu/fitsm