



# EGI POSITION PAPER FOR A FEDERATED EUROPEAN CLOUD

**Authors:** Michel Drescher, EGI.eu,  
Tiziana Ferrari (EGI.eu)  
David Wallom (U.Oxford)

**Date:** 4 December 2013

**Document status:** FINAL

**Document link:** <https://documents.egi.eu/document/2050>

## Abstract

This position paper discusses the concept of a European Federated Cloud. This expands from the initial service offering from the EGI's Federation IaaS Cloud Infrastructure Platform. This builds on the foundations offered by the properties and services of EGI's Core Infrastructure Platform supporting a set of infrastructure independent services to facilitate federation of resources. A rich portfolio of Cloud services is expected to emerge over time, which will be further facilitated through directed actions coming from within EGI, local, national and European funding streams.

The document concludes with a brief overview of the value proposition for the two principal target audiences of this document; Cloud service providers on one hand, and Cloud service consumers on the other.

## COPYRIGHT NOTICE

Copyright © Members of the EGI-InSPIRE Collaboration, 2010. See [www.egi.eu](http://www.egi.eu) for details of the EGI-InSPIRE project and the collaboration. EGI-InSPIRE ("European Grid Initiative: Integrated Sustainable Pan-European Infrastructure for Researchers in Europe") is a project co-funded by the European Commission as an Integrated Infrastructure Initiative within the 7th Framework Programme. EGI-InSPIRE began in May 2010 and will run for 4 years. This work is licensed under the Creative Commons Attribution-Noncommercial 3.0 License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc/3.0/> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, and USA. The work must be attributed by attaching the following reference to the copied elements: "Copyright © Members of the EGI-InSPIRE Collaboration, 2010. See [www.egi.eu](http://www.egi.eu) for details of the EGI-InSPIRE project and the collaboration". Using this document in a way and/or for purposes not foreseen in the license, requires the prior written



permission of the copyright holders. The information contained in this document represents the views of the copyright holders as of the date such views are published.



## Table of Contents

<b>Abstract</b> .....	<b>1</b>
<b>1 Introduction</b> .....	<b>4</b>
1.1 Challenges for the ERA and the Digital Agenda .....	4
<b>2 The EGI federated Cloud</b> .....	<b>5</b>
2.1 Platform architecture .....	5
2.2 Service tiers .....	7
2.2.1 <i>Reliable Infrastructure Cloud</i> .....	7
2.2.2 <i>General-purpose platform services</i> .....	7
2.2.3 <i>Professional research infrastructure platforms</i> .....	8
2.2.4 <i>Zero ICT infrastructures</i> .....	9
2.3 Business models.....	9
2.3.1 <i>Customers for federated Cloud services</i> .....	9
2.3.2 <i>Resource management and pricing</i> .....	10
2.3.3 <i>Service Level Agreements</i> .....	11
<b>3 Value proposition</b> .....	<b>12</b>

## 1 Introduction

The EGI (European Grid Infrastructure) is a federation of shared computing, storage and data resources from the public sector through national and intergovernmental resource providers and a fledgling set of commercial resource providers. These deliver integrated and secure distributed computing, storage and data services to European researchers and their international collaborators. Current production services support large scale data processing mainly through the capability of running computational jobs over a distributed infrastructure and offering data management services for storing, retrieving, accessing, replicating files in a secure manner.

### *1.1 Challenges for the ERA and the Digital Agenda*

The current production services are key to support certain computing models and have demonstrated their value for a number of scientific use cases. Nonetheless, these computing models reflect needs for some communities that have largely overlapping e-Infrastructure requirements through a limited common application design pattern. To increase the overall opportunities for sustainability, EGI will diversify its support for a much wider set of research communities in Europe, and evolve into a major infrastructure stakeholder for the European Research Area (ERA).

Supporting a much wider set of research communities unequivocally results in needing to accommodate a much more diverse set of e-Infrastructure requirements computing models and application design patterns. Naturally, such a strategic change of scope will require the production infrastructure management to respond with a strategy to cope with these new requirements.

With the technological evolution of the last years, virtualisation has become a commodity, and support at the hardware level has removed performance bottlenecks in many of the most of the common use cases. These were contributing factors in the development of cloud computing where infrastructure providers built on-demand and self-service provisioning IT infrastructures together with a number of complimentary services and management tools.

This paradigm shift represents a great opportunity for EGI to address the ICT needs of different research communities and to gradually evolve the overall infrastructure into a modular set of platforms that, at different levels of abstraction, can support the deployment of software stacks from various communities.

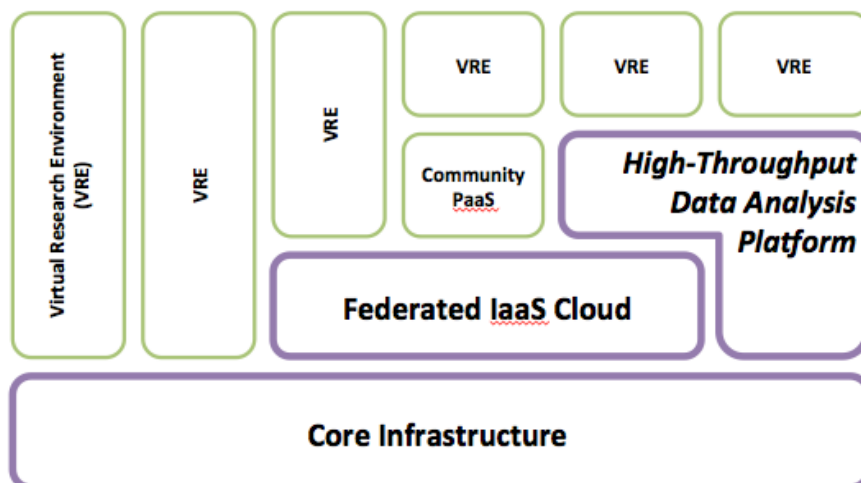
## 2 The EGI federated Cloud

EGI's Federated IaaS Cloud model is built on the concept of service orientation, which is core to the Cloud computing paradigm. This means that the federation develops, defines and enforces the delivery of a well-defined set of services, and packages them into a delivery model, where each participant fulfils a clearly defined role within the federation. These services are defined around capabilities, not technologies: Technical services may deliver the functionality comprising more than one capability, but the federation neither specifies nor mandates which technology is to be used to deliver this service. Instead, the federation describes the necessary service interface functionality, and seeks satisfy this with publicly defined open specification. Where such specifications match the required functionality, the federation mandates the support for such specification, based on consensus building among the members.

As a consequence, members of the federation are free to choose whichever technical implementation they prefer for the type of service they are providing as part of the federation, for as long as the service exposes the mandated interface.

### 2.1 Platform architecture

The foundation of EGI's federated IaaS Cloud infrastructure is provided by the technical architecture for EGI's production infrastructure that is modelled around software platforms.



**Figure 1: Platform-oriented architecture of EGI's production infrastructure<sup>1</sup>**

Of the three pillars of EGI's strategy for H2020, the **operational infrastructure** provides the technical services that implement and provide the functionality necessary to address the needs of EGI's user communities.

EGI as a foundation does not own the physical resources on which the technical services are deployed – but (a subset of) the members of EGI do. This is consequently reflected in the technical architecture as follows.

<sup>1</sup> Collaboration services are not shown in this diagram.

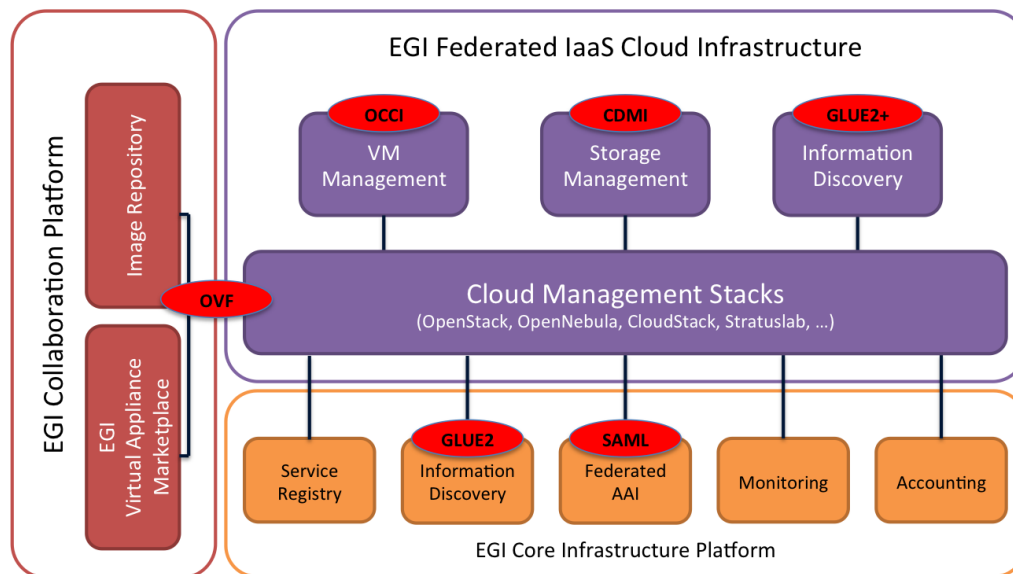
The **Core Infrastructure Platform** comprises of technical services that are necessary to build a federation of services and other platforms. It includes key subsystems required to operate a federated infrastructure, such as:

- Federated AAI based on PKI and X.509 certificates
- Monitoring of the entire infrastructure
- Accounting of consumed resources
- Information discovery
- Service registry for deployed services in EGI

Any *other* platform EGI accepts for deployment and integration in the production infrastructure must integrate with the Core Infrastructure Platform services. A number of Grid middleware platforms are already in the process of integrating with the Core Infrastructure Platform, or have already accomplished this:

- EMI-3 HTC Grid platform
- IGE-2 Grid platform
- UNICORE HPC Grid platform
- QCG multi-scaling & reservation platform

The **High-Throughput Data Analysis Platform** is EGI's successor of the EMI-3 platform developed, integrated and packaged by the European Middleware Initiative (EMI) project. It largely comprises of HTC Grid middleware services such as CREAM, ARC, dCache, etc.



**Figure 2: Federated IaaS Cloud Infrastructure Platform**

The EGI **Federated IaaS Cloud** is the third infrastructure platform that EGI is specifying in terms of scope, capabilities, included components, and business models around it. Complying with EGI's requirement of any other platform to integrate with the Core Infrastructure platform, EGI's federated IaaS Cloud inherits AAI, Messaging, Monitoring and Accounting capabilities, and IaaS Cloud services are properly registered in EGI's central service registry. Currently, the federated IaaS Cloud infrastructure comprises of the following functionality:

- Cloud compute resources (through OCCI endpoints)



- Cloud storage resources (through CDMI endpoints)
- VM Image distribution infrastructure
- Cloud-special information discovery
- EGI's Virtual Appliance Marketplace<sup>2</sup>

Together<sup>3</sup>, these three platforms provide the foundation for a production infrastructure that can support many different and diverse user communities in the ERA.

## 2.2 Service tiers

As part of EGI's Solution catalogue, the EGI Federated IaaS Cloud solution comprises a number of services, each with an associated service delivery model as follows.

### 2.2.1 Reliable Infrastructure Cloud

This service provides infrastructure Cloud resources (i.e. Cloud compute & storage resources) to agreed service levels for use by any EGI user community.

This service delivered as a collaboration of federated EGI members, EGI.eu and other service providers as follows:

#### EGI.eu

- Service level monitoring
- Resource usage accounting
- Federated AAI coordination
- Service discovery
- VM Image distribution network
- EGI Virtual Appliance Marketplace

#### EGI Resource Providers:

- IaaS Cloud compute resources
- IaaS Cloud storage resources
- Information discovery

**Other provisions** (may also be individually delivered by EGI.eu or EGI resource providers):

- Service-brokering and scheduling
- User community management (VO management)
- Virtual Appliance image repository

### 2.2.2 General-purpose platform services

This service comprises of a number of individual software systems provided as a reliable service, to agreed service levels for use by any EGI user community.

Software systems that fall into this category are systems that EGI might package and integrate into one of more professional research infrastructure platforms that are deployed on top of EGIs federated IaaS Cloud. Additionally these also have a market to be delivered as a service in a stand-alone fashion.

---

<sup>2</sup> The VA Marketplace is an extension to the EGI Application database inheriting all collaborative features while allowing managing Virtual Appliance metadata and integration with the VM Image distribution infrastructure

<sup>3</sup> Including where required collaboration services that are part of the EGI Collaboration Platform



Examples for such services might be:

- Generic, general-purpose databases (PostgreSQL, MySQL, MongoDB, etc)
- Hadoop compute cells as a service
- VM Image housekeeping and cleanup service
- Persistence Identifier resolver/replica

While the vast majority of services falling into this category will be deployed as platform services, EGI expects a number of software systems to be deployed and offered in a Software as a Service (SaaS) model.

**EGI Resource Providers** may provide services that fall into this category.

### 2.2.3 Professional research infrastructure platforms

Research infrastructure platforms comprise of a well-defined, integrated set of software systems that together provide the necessary and required functionality to agreed service levels for use by any EGI user community. In fact, just like the general-purpose platform services described above, this service is a collection of specific research infrastructure platforms that are individually provided as a service, to individually agreed service levels.

Such platforms are expected to make provisions for self-service and elasticity functionality as typical Cloud service characteristics, both on the software system level (e.g. self-service configuration interface for scaling up and down certain software system properties) and the platform composition level (e.g. adding or removing platform components from the consumed service).

For services that fall under this category, it is expected that formal or informal **EGI special interest groups for Research Infrastructure (SIG RIs)** will form around the service provisions for individual research infrastructure platforms. The specific service delivery composition highly depends on the architecture of the research infrastructure platform, and will have to be defined on a case-by-case basis. However, such SIG RIs are expected to entirely emerge from within the EGI federation.

Candidate research infrastructure platforms begin to emerge through EGI's numerous collaborations with many other publicly funded projects, such as WeNMR, Lifewatch, DCH-RP, EISCAT-3D, SCIDIP-ES, APARSEN, Peachnote, BioVel, EUBrazilOpenBio, among others. The following examples provide an indication of the type of possible platforms:

#### **Secure data storage Platform:**

Stored data is protected from unauthorized access; even system administrator access is prohibited or, where absolute protection is either impossible or not necessary, adequately logged for retroactive privacy protection audits. Key components of such a platform are: (i) key management, (ii) server side encryption, (iii) secure data deletion, and (iv) ACL based data sharing.

Such secure data storages are designed and offered as a *safe for storing (personal) confidential data*.

#### **Data dissemination/Open Data Platform:**

Cloud storage provides an infrastructure to collect and disseminate scientific data. Data ingestion and management is curated, data access is inherently open (anonymous) or tracked (e.g. through accepting social identities for data access).





Key components of such a platform might be: (i) Generic Cloud Object Storage, (ii) Custom Data Access Control Lists, (iii) SaaS-delivered data catalogues, (iv) Federated (social) AAI for data access, (v) Data access tracking service

### **Data preservation platform:**

A platform for long-term data preservation: Born-digital or digitized data is prepared on-site for ingestion in the off-site preservation platform provided by EGI. The preservation platform may include automatic and automated data integrity checking, data format conversion services, file format standard compliancy checks, data redundancy, disaster management through cross-continent storage dispersion, and transparent storage technology transition. Such a preservation service should be considered fit for external certification against recognised digital archive procedures and best practices.

### **2.2.4 Zero ICT infrastructures**

This service comprises of a number of research infrastructure platforms that are provided as a reliable service, to agreed service levels for use by any EGI user community.

The difference to professional research infrastructure platforms services lies in the deployment model of the included software components and stacks. While the former service exposes programming and scripting interfaces to the end user, services provided in this tier are designed so that Web browsers or mobile devices are the only *required* ICT equipment for the end user. This does not exclude such platforms to also expose APIs for the convenience of the user, but these must be *optional*.

The scope of such services is not limited to electronic services. Zero ICT infrastructure services may also entail software for managing research infrastructure equipment, and facilities (e.g. shift management, facilities management).

Very similar to research infrastructure platform services mentioned above, Zero ICT infrastructure services will likely see EGI SIG Zero ICT (EGI SIG ZIT) to emerge around the service provisions of such services.

## **2.3 Business models**

The business models around the EGI federated Cloud infrastructure are categorised as follows:

- Customer scope and segments
- Resource management and pricing
- Service level agreements

### **2.3.1 Customers for federated Cloud services**

The current scope of EGI as a federation of resources and infrastructure services clearly lies around providing services for and to members of EGI as a federation. This constraint is in place by practice, not by statutes. The EGI federation terms of reference and membership rules are not limited to academic resource providers; by no means SME and large industrial entities are excluded from any form of membership in EGI.

EGI currently comprises members coming almost entirely from the academic sector (represented as NGIs) or European research organisations (EIROs such as CERN and EMBL). Such members own the resources on which the EGI production infrastructure is deployed. Most if not all of these resources were purchased using funds with attached usage constrains, for example to make available at least 80% of all purchased resources for the High Energy Physics research domain supported by this

resource centre. This is the root cause of the natural, practical scope of EGI currently serving its members in a cooperative solidary business model.

However, EGI will have to decide whether to award membership to resource providers that do not have such resource allocation constraints, as this will have a fundamental impact on the targeted market segments: EGI then may be in the position of serving customers that are not in any way affiliated with members of EGI. This has a number of implications that need to be discussed and decided upon on the EGI Council level.

### **2.3.2 Resource management and pricing**

Any provision of services incurs a cost. There follows a breakdown of the currently envisaged EGI business model for Cloud resources. We provide an overview of current pricing models, for which individual Cloud resource and service providers may indicate their individual cost tags. Several of the later models are similar to those that are used by public cloud providers. Please note that, by concept, all pricing models are available on all defined service tiers in EGI (see above).

#### **1. Free at the point of use (FAUP)**

Resource and service consumption will be accounted, aggregated, and made available to the service customer. Costs associated to delivering resources and services on this plan will be indirectly recovered, through external local, national or European funding (e.g. Project proposals).

#### **2. Try before you buy (TBB)**

A pre-defined collection of resources or services will be made available for prospective customers to explore the capabilities of the offered service, with some limitation applied to either the volume of resource or the time for which that resource is available. There is a large overlap between FAUP and TBB priced resources, in that these are often the preferred models in funded project settings (e.g. national or European funded projects).

TBB priced resources are often used to facilitate the conversion rate from non-paying customers to paying customers.

#### **3. Pay as you go (PAYG)**

This pricing model makes no connection between provision for resources and consumption. This pricing model moves all responsibility onto consumers who pay only for those resources that they consume; collection of moneys can happen either by invoicing the customer, or by customers depositing funds with the resource and/or service provider in a prepayment scheme, also, standing order payment schemes may be eligible for cost recovery.

#### **4. Wholesale purchases**

Wholesale purchases link resource and service provisions to the actual purchase. This means that wholesale purchases are effectively a pre-financing scheme for resource provisions, allowing the provider to better plan service continuity and capacity. The consequence is a lower price when compared to PAYG plans (which, by virtue, are post-financing resources through cost recovery).

Wholesale purchases follow a model similar to that of a bank account, i.e. customers are allowed to over-consume resources up to a certain allowance. Similar to agreed current account overdrafts, resource consumption in excess of the wholesale purchase will be costed under the PAYG pricing plan.

## 5. Reserved resources

This pricing plan provides customers with exclusive access to the purchased resources, for the agreed amount of time. This model is *not* a flat rate model, in that resources are indeed reserved exclusively for the respective customer, irrespective of the actual use of the resources.

Resource and service consumption will be allowed up to the reserved amount, but no further. Alternatively, this pricing plan may follow the overdraft concept as well, in that excess consumption will be priced according to the PAYG plan.

### 2.3.3 Service Level Agreements

Another aspect of the EGI business model deals with service levels for provided services. Within the remit of EGI and services provisioned and consumed under the auspices of EGI, all services will be delivered to clear, pre-formulated Service Level Agreements.

Currently, EGI is developing the following Service Level Agreement tiers:

- **Best Effort**  
Best effort service levels are available for all service tiers (see above) and all Cloud service providers. Best effort service levels make no promise or guarantee for any service level.
- **Basic**  
Basic Service Level Agreements are available for all EGI service tiers, and all Cloud service providers. However, Basic Service Levels are available only on pricing plans PAYG or higher.  
The exact definition of Basic depends on the service tier in EGI.
- **Custom SLAs tailored to community requirements**  
These comprise of advanced SLAs subject to specific requirements or research infrastructures and innovation projects. Also, not all Cloud service providers within EGI may provide such advanced SLAs.  
Custom SLAs are available for all EGI service tiers.

### 3 Value proposition

The EGI Federated Cloud makes a compelling case for a multi-tier service platform EGI can capitalise on in its support of user communities coming from the publicly funded research sector and, in the future, with new types of resource provider joining, other communities stemming from the commercial or public sector in Europe.

The European Federated Cloud, driven by and through EGI, has many benefits for its stakeholders:

#### **Cloud service providers**

- Increase utilisation of IaaS Cloud resources through federation membership,
- Increased market visibility and penetration through service delivery in the federated Cloud,
- Shared cost and effort of first line support for users,
- Shared and common CSIRT processes,
- Future development & coordination,
- A technically challenging certification provides a shared competitive edge over non-federation service providers,
- Flexible Cloud service tiers for value added service provisioning on top of the EGI IaaS Cloud federation foundation
- Flexible business model framework
- Focused user community support through shared interest in various special interest groups

#### **Cloud service consumers**

- Simplified Cloud resource access management through consistent access and AA interfaces
- Service providers are comparable through centralised resource reliability monitoring across providers
- Fair service provisions through pre-fabricated Service Level Agreements across service providers
- Easy self-service provisioning through advanced services on the IaaS service tier
- Flexible multi-tier service platform provides a rich European service portfolio to choose from
- Common Code of Conduct agreement across service providers in the EGI federation