



EGI-InSPIRE

EGI PLATFORMS ROADMAP

EU MILESTONE: MS514

Document identifier:	EGI-InSPIRE-MS514-FINAL
Date:	24/04/2013
Activity:	SA2
Lead Partner:	EGI.eu
Document Status:	FINAL
Dissemination Level:	PUBLIC
Document Link:	https://documents.egi.eu/document/1624

Abstract

This document introduces the platform-driven service delivery model to the EGI community. It defines the term platform and how IT platforms fit into the current and emerging EGI ecosystem. After providing an overview of the EGI platform architecture, the first part of the document describes the different platforms in more detail. The second part of the document provides a roadmap on the adoption of the platform-based architecture and the service delivery model used by the European Grid Infrastructure.



I. COPYRIGHT NOTICE

Copyright © Members of the EGI-InSPIRE Collaboration, 2010. See 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 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.

II. DELIVERY SLIP

	Name	Partner/Activity	Date
From	Michel Drescher	EGI.eu	
Reviewed by	Ludek Matyska	Elixir	14/4/2013
Moderated	Ignacio Blanquer		15/4/2013
Approved by	AMB & PMB		24/4/2013

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	20-03-2013	Initial version.	Michel Drescher, EGI.eu
2	21-03-2013	First draft complete.	Michel Drescher, EGI.eu
4	28-03-2013	Incorporated internal review comments	Michel Drescher, EGI.eu
5	02-04-2013	Draft ready for external review	Michel Drescher, EGI.eu
6	15-04-2013	Incorporate Reviewer comments	Michel Drescher, EGI.eu
7	17-04-2013	Draft ready for AMB/PMB review	Michel Drescher, EGI.eu
8	18-04-2013	Review AMB changes, image formatting	Michel Drescher, EGI.eu

IV. APPLICATION AREA

This document is a formal deliverable for the European Commission, applicable to all members of the EGI-InSPIRE project, beneficiaries and Joint Research Unit members, as well as its collaborating projects.

V. DOCUMENT AMENDMENT PROCEDURE

Amendments, comments and suggestions should be sent to the authors. The procedures documented in the EGI-InSPIRE “Document Management Procedure” will be followed:

<https://wiki.egi.eu/wiki/Procedures>

VI. TERMINOLOGY

A complete project glossary is provided at the following page: <http://www.egi.eu/about/glossary/>. Additional definitions of terms may be found in the ITIL 2011 Glossary [R 11] and the EGI Technology Glossary [R 12].

The following table provides a set of terms that are used in this document.

Term	Description
EGI Platform model	The EGI Platform model refers to business models that may emerge by utilising any of the IT platforms that are described in the EGI Platform architecture.
EGI Platform architecture	Describes how the individual platforms (see below) are embedded in the EGI ecosystem, and how they are technically integrated with the current EGI production infrastructure.
EGI Core Infrastructure Platform	The EGI Core Infrastructure Platform consists of services that are necessary to establish and operate a federated, distributed computing and data infrastructure (DCDI). This platform does <i>not</i> include services that provide access to computing or data resources. In EGI, any other platform that is deployed and operated by EGI <i>must</i> integrate with the EGI Core Infrastructure Platform.
EGI Cloud Infrastructure Platform	The EGI Cloud Infrastructure Platform provides a federation of IaaS Clouds deployed in the EGI ecosystem. Integrated with the EGI Core Infrastructure Platform, it provides consistent and standardised access to virtualised compute, storage and networking resources.
EGI Collaboration Platform	The EGI Collaboration Platform provides IT Infrastructure and Services that facilitate collaboration between Research Communities without being a core infrastructure service for Research Communities.
Community Platform	Community Platforms (there may be more than one) consist of services that are specific to or preferred by the respective community. Community Platforms may be deployed on top of the EGI Core Infrastructure Platform, or on top of the EGI Cloud Infrastructure Platform.



VII. PROJECT SUMMARY

To support science and innovation, a lasting operational model for e-Science is needed – both for coordinating the infrastructure and for delivering integrated services that cross national borders.

The EGI-InSPIRE project will support the transition from a project-based system to a sustainable pan-European e-Infrastructure, by supporting ‘grids’ of high-performance computing (HPC) and high-throughput computing (HTC) resources. EGI-InSPIRE will also be ideally placed to integrate new Distributed Computing Infrastructures (DCIs) such as clouds, supercomputing networks and desktop grids, to benefit user communities within the European Research Area.

EGI-InSPIRE will collect user requirements and provide support for the current and potential new user communities, for example within the ESFRI projects. Additional support will also be given to the current heavy users of the infrastructure, such as high energy physics, computational chemistry and life sciences, as they move their critical services and tools from a centralised support model to one driven by their own individual communities.

The objectives of the project are:

1. The continued operation and expansion of today’s production infrastructure by transitioning to a governance model and operational infrastructure that can be increasingly sustained outside of specific project funding.
2. The continued support of researchers within Europe and their international collaborators that are using the current production infrastructure.
3. The support for current heavy users of the infrastructure in earth science, astronomy and astrophysics, fusion, computational chemistry and materials science technology, life sciences and high energy physics as they move to sustainable support models for their own communities.
4. Interfaces that expand access to new user communities including new potential heavy users of the infrastructure from the ESFRI projects.
5. Mechanisms to integrate existing infrastructure providers in Europe and around the world into the production infrastructure, so as to provide transparent access to all authorised users.
6. Establish processes and procedures to allow the integration of new DCI technologies (e.g. clouds, volunteer desktop grids) and heterogeneous resources (e.g. HTC and HPC) into a seamless production infrastructure as they mature and demonstrate value to the EGI community.

The EGI community is a federation of independent national and community resource providers, whose resources support specific research communities and international collaborators both within Europe and worldwide. EGI.eu, coordinator of EGI-InSPIRE, brings together partner institutions established within the community to provide a set of essential human and technical services that enable secure integrated access to distributed resources on behalf of the community.

The production infrastructure supports Virtual Research Communities (VRCs) – structured international user communities – that are grouped into specific research domains. VRCs are formally represented within EGI at both a technical and strategic level.



VIII. EXECUTIVE SUMMARY

This document provides an update of the initial EGI Platform Roadmap [R 3] which, for the first time, grouped the services offered by EGI's Resource Infrastructure Providers, based around the consumers of these services and the value they offered. This was a necessary first step in being able to discuss how to grow the usage of these services and broaden adoption.

The EGI Platform Architecture has matured into its current form as described in the first part of this document. The EGI Core Infrastructure Platform is scoped around the mechanisms and tools needed for federating Community Platforms, and a newly established EGI Cloud Infrastructure Platform focuses on Cloud Computing as a means to provide federated consistent access to DCIs across Europe. Consequently, the EGI Cloud Infrastructure Platform integrates with the EGI Core Infrastructure Platform. The services provided within the Unified Middleware Distribution (UMD) are an example of another Community Platform that is integrated with EGI's Core Infrastructure Platform.

With the increasing maturity of the technical platforms, the stakeholders in the EGI community become better defined, and the roles they may play in different deployment and integration scenarios become clearer than before.

What was considered a concept before is now swiftly becoming an agreed architecture of the EGI production infrastructure, both (partially) driven *and* validated by the EGI Cloud Infrastructure Platform. With existing and new services finding their place in this architecture the composition of each EGI Platform allows for formalisation and persistent documentation of the IT landscape in EGI. The first steps have been made in this document, by providing the general architecture and integration between components across platforms. The next steps will have to be further decomposing the architecture into finer details.

However, equally important is to further work on the actual IT service development around the technical/software services that comprise the EGI production infrastructure. These IT services will help EGI in offering consumption models and support to its customer base at large. Although part of the second pillar of the EGI Strategic Plan, these services are closely tied into the technical infrastructure evolution comprising the first pillar of EGI's strategy to warrant brief discussion in this document. Taking the platform architecture's abstraction level concept and applying it to the IT service development, EGI will utilise the TCB and the newly formed URT to strike the right balance between business relationship formality and flexibility in community building.

Consequently, the EGI Platforms Roadmap will *also* focus on further development of the IT services based on the technical components described in this document.



TABLE OF CONTENTS

1	INTRODUCTION	7
2	EGI PLATFORM ARCHITECTURE: AN UPDATE	8
2.1	Stakeholders and Actors	9
2.2	Roles and responsibilities in the EGI ecosystem	10
2.3	EGI Core Infrastructure Platform	12
2.4	EGI Cloud Infrastructure Platform	13
2.5	EGI Collaboration Platform	15
2.6	Community Platforms	16
3	EGI PLATFORM ROADMAP: SUPPORT MODELS AND PROCESSES ...	18
3.1	Maturing EGI Platforms	19
3.2	Developing IT Services and Service Offerings	19
3.3	Engaging with Technology Providers	19
3.4	Implementing new support structures	19
3.5	Define IT services provided by Technology Providers	20
3.6	Preparing EGI Platforms beyond EGI-InSPIRE	20
4	CONCLUSION	21
5	REFERENCES	22

TABLE OF FIGURES

Figure 1:	The EGI Platform Architecture	8
Figure 2:	Typical distribution of roles in the EGI Grid ecosystem	10
Figure 3:	Relationships around Community Platforms deployed onto EGI's Core Infrastructure	11
Figure 4:	Relationships around Community Platforms deployed in EGI's Cloud Infrastructure Platform	12
Figure 5:	Components of the EGI Core Infrastructure Platform	13
Figure 6:	Architecture of the EGI Cloud Infrastructure Platform	14
Figure 7:	Current realisations of the abstract Cloud Management stack component	15
Figure 8:	EGI Collaboration services capture social and technical collaboration needs	15
Figure 9:	A typical structure of IT Service offerings	18

TABLE OF TABLES

Table 1:	Stakeholders and Actors in the EGI Platform ecosystem	9
Table 2:	Provisional list of Technology Providers	17



1 INTRODUCTION

The EGI Strategy [R 1] outlines how EGI is gradually yet continuously aligning itself with the EC's vision of H2020. Consequently the EGI Strategy is built around implementing three pillars of services to the European Research Area (ERA); (1) Provide a pan-European e-Infrastructure, (2) Provide Community and Coordination services, and (3) Foster the development and support for Virtual Research Environments in Europe. Within this strategy, the technical infrastructure forms the first pillar of EGI's support for the researchers and research communities that make up the ERA. Part of this overarching goal is to define the scope for the EGI platforms and mature their respective technical architecture into a model that allows access to and integration with these platforms as a marketable asset to research communities within the EGI ecosystem.

EGI will continue supporting its current customers, the Heavy User Communities (HUC) comprising of the high-energy physics, astrophysics, structural biology and other traditional communities in the EGI ecosystem. At the same time however EGI is positioning itself to support much more diverse research communities, smaller research teams; generally what is referred to as "the long tail of science". These objectives, goals and activities are described in varying levels of detail and aspects in the EGI Strategy [R 1], EGI Technical Roadmap [R 2], the first edition of the EGI Platform Roadmap [R 3] and other documents.

This document is the second edition of the EGI Platform Roadmap. It is not written as a true edition, i.e. amending the complete contents where required. It rather provides an *update* to the previous edition, referring to it where required. Therefore, to fully understand this document, the reader is invited make themselves familiar with the EGI Platform Roadmap, first edition [R 3] before continuing with this document.

Hence this document is fundamentally structured around two central sections as follows:

Section 2 provides an update on the state of the art of the EGI Platform Architecture since the first edition of this document, briefly recapturing each EGI Platform's scope. This includes an update on the stakeholder/actor model around these platforms, reflecting the discussions and need for guidance in the EGI community. This section concludes with brief updates on the architectural changes in each EGI platform and a summary of considered Community Platforms in EGI.

Section 3 presents the roadmap for the EGI Platform Architecture until April 2014, which in itself forms a major hallmark in the evolution of the activity itself. The roadmap is intentionally kept on a high level, providing guidance and direction to more detailed roadmap descriptions in the second edition of the EGI Technical Roadmap [R 4].

This document concludes with section 4 summarising the progress over the last year while briefly rationalising the direction of the EGI Platform Roadmap for the coming year.

2 EGI PLATFORM ARCHITECTURE: AN UPDATE

The concepts of the EGI Platform architecture were extensively described in [R 3]. From this starting point the EGI Platform architecture gradually evolved and matured into what it is described in this updated document (see Figure 1). Since then, the EGI Platform Architecture was the topic of many discussions, both online and offline within the EGI community. In retrospect, these discussions were all focused on the sustainability of EGI. Both driving and reflecting these discussions, the EGI Platform Architecture responds to this most urgent concern in the EGI community on a technical level:

Which is, for the EGI ecosystem, the most suitable technical architecture in order to achieve sustainability?

Focusing on technical sustainability alone may be sufficient only in scenarios that include a stable business and revenue model. EGI, however, is also confronted with fundamental changes in the business and revenue models. The future of EC-funded projects for Grid middleware development is very unclear once the IGE and EMI projects have ended, as it is the case for operational activities within EGI beyond the EGI-InSPIRE project. Influenced by and exposed to the discussions around the EGI business model, the EGI Platform Architecture responds to these challenges with two fundamental changes:

- Amend the EGI production infrastructure with generic infrastructure capabilities – the EGI Cloud Infrastructure Platform
- Re-group parts of the existing production infrastructure into a marketable IT service contributing to cost recovery for operating included services

Choosing a platform approach allows the EGI community, particularly the participating NGIs, to develop value added services on top of any of the EGI Platforms. The most visible outcome of this change is in fact the EGI Cloud Infrastructure Platform itself, both driving and validating the EGI Platform Architecture through day-to-day discussions within the EGI Federated Clouds Task in SA2.

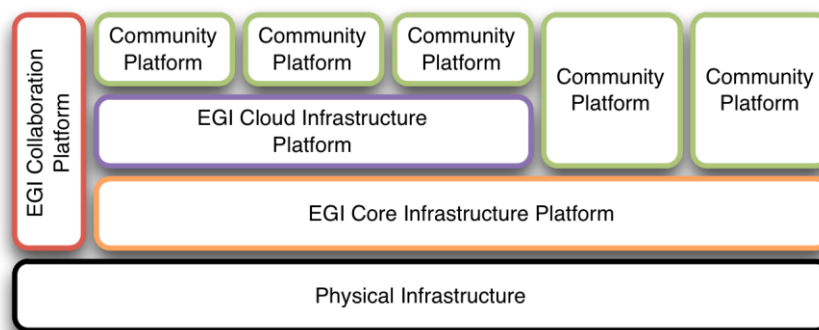


Figure 1: The EGI Platform Architecture

The Core Infrastructure Platform as described in MS510 [R 3] is now organised in two distinct platforms as described below: The EGI Core Infrastructure Platform (CIP) and the EGI Cloud Infrastructure Platform (CLIP).

The **EGI Core Infrastructure Platform** is now re-scoped to include only services that are necessary to establish and operate a federated distributed computing infrastructure. This implies that, unlike before, there are no core resource access services included in the EGI Core Infrastructure Platform.



The underlying strategic decision was introduced in the EGI Strategic Plan [R 1] and the accompanying Technical Roadmap [R 2]. From there on this concept was further evolved and discussed with major stakeholders, e.g. at the EGI Technical Forum [R 5] and the Evolving EGI workshop in January 2013 [R 6]. This division allows the EGI Core Infrastructure Platform to be offered as a service to other e-Infrastructures and Research Infrastructures in Europe and around the world, without imposing a certain resource access and usage paradigm to the partnering e-Infrastructures or Research Infrastructure.

To provide generic, consistent and flexible access to EGI resources, EGI initiated a strategic activity to establish a federation of locally deployed IaaS Clouds. The **EGI Cloud Infrastructure Platform** directly supports EGI’s strategic alignment with the European Commission’s Horizon 2020 strategy [R 7]. While EGI will continue to support and maintain its existing relationships with research communities (such as WLCG), it will provide the Cloud Infrastructure Platform in support of new research communities who wish to deploy their own Virtual Research Environment on the resources that they are able to access. This enables less structured research collaborations and individual researchers stemming from the so-called “long tail of science” to easily access and configure EGI’s resources.

Complementing the two previously outlined platforms, the **EGI Collaboration Platform** will continue to offer services that allow collaboration across research communities and their domain-specific community platforms. It facilitates synergies between research communities by encapsulating services that are common across multiple communities and are not critical to the operation of the EGI production infrastructure, therefore are outside of the Core and Cloud platforms. The EGI Collaboration Platform complements the other platforms and contributes to their efficient use.

During the last year of evolution of the EGI Platform architecture, a number of services were relocated from one platform to another. The following sections will describe the individual EGI platforms in greater detail.

2.1 Stakeholders and Actors

Since the first edition of this document, the actors and stakeholders have slightly changed into the model described in this section.

Stakeholder	Actor Name	Actor Verb
Resource Infrastructure Provider	Platform Owner	Own
European Grid Infrastructure	Platform Contributor	Contribute
Technology Provider	Platform Packager	Package
Platform Integrator	Platform Deployer	Deploy
Research Community	Platform Operator	Operate
	Platform Users	Use

Table 1: Stakeholders and Actors in the EGI Platform ecosystem

Table 1 provides an overview of the current stakeholders and actors; the roles have not changed and stay the same as described in [R 3]; except for the following changes:

The **European Grid Infrastructure**, as a federation of Resource Infrastructure Providers (mostly NGIs), captures EGI Federation members' vested interest in pan-European services. In the context of this document, this comprises central technical services that members of the EGI federation outsource to one of their members, coordinated through EGI.eu.

The **Platform Contributor** is a new actor capturing the role of Technology Providers that develop and maintain software that is included in any platform in the EGI ecosystem. The term *contributor* does not imply an active role in the development and maintenance of any given platform. In fact, it is often the case that contributing is a *passive* role in this ecosystem, particularly for platforms that depend on Open Source software and projects.

The **Platform Operator** is now considered an actor instead a stakeholder, reflecting the ecosystem reality more accurately than before – Platform Operators are overwhelmingly if not exclusively recruited from within existing EGI ecosystem stakeholders.

These changes have a number of implications on the applicability of actors and stakeholders in the deployment scenarios for Community Platforms. These are illustrated in the following subsection, including the EGI Cloud Infrastructure platform and the new stakeholder and actor distribution.

2.2 Roles and responsibilities in the EGI ecosystem

Next to the EGI Platforms (currently, EGI Core Infrastructure, EGI Cloud Infrastructure and EGI Collaboration Platforms) there are a number of Community Platforms that need to be managed in an appropriate way. The Actor roles described in the previous section apply to all stakeholders, although differing across platforms. For example, Resource Infrastructure Providers will deploy and operate Community Platforms that directly integrate with the EGI Core Infrastructure Platform (Figure 3). However, Research Communities themselves typically assume these roles for their Community Platforms deployed on top of the EGI Cloud Infrastructure Platform (Figure 4).

The following illustrations will focus on the relationships and interactions around and between the Community Platforms and the underlying EGI Core Infrastructure and EGI Cloud Infrastructure Platforms. This is not to disregard the EGI Collaboration Platform as less important; it is not included in the following section for the sake of clarity and brevity since the roles and relationships are identical to those of the EGI Core Infrastructure Platform.

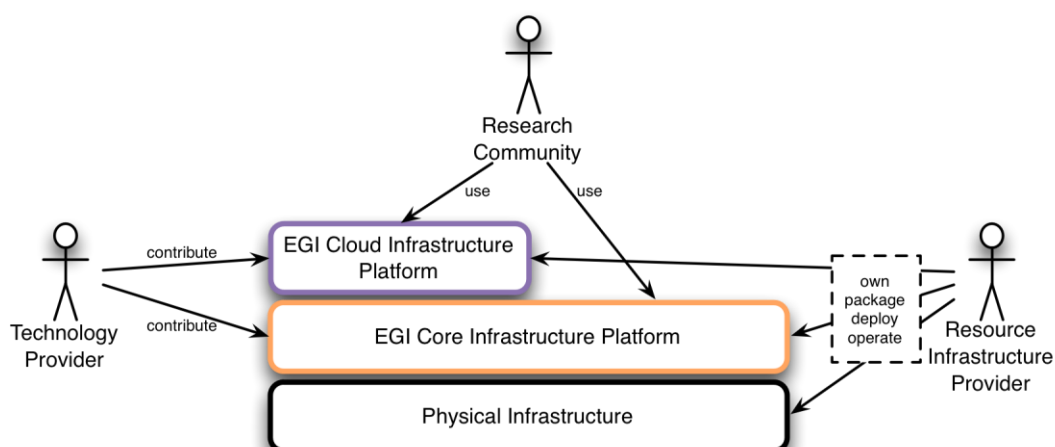


Figure 2: Typical distribution of roles in the EGI Grid ecosystem

Figure 2 illustrates which stakeholder in the EGI ecosystem takes on which Actor role by indicating them as verbs qualifiers on the association arrows. Research Communities are the main customers of the EGI Platforms (Core and Cloud) and use them through their own Community Platforms (not shown).

While the individual Resource Infrastructure Providers deploy and operate both the EGI Core and Cloud Infrastructure Platforms as part of the EGI federation, EGI acts as their own Platform Integrator through a number of services and activities that are coordinated on the European level¹, supported by the EGI-InSPIRE project. The Resource Infrastructure Providers literally own the hardware resources that are exposed and federated through the illustrated platforms; however it is the federation of Resource Infrastructure Providers (EGI) that exerts shared ownership on the EGI Core, Cloud and Collaboration Platforms.

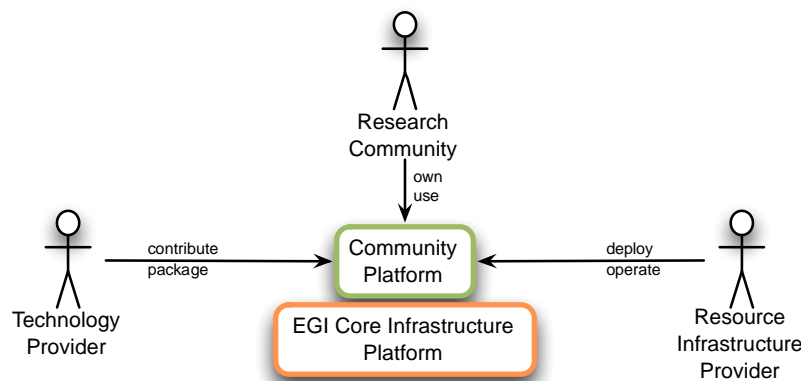


Figure 3: Relationships around Community Platforms deployed onto EGI’s Core Infrastructure.

Figure 3 describes the roles stakeholders assume for Community Platforms that are deployed directly on the hardware resources and integrated with the EGI Core Infrastructure Platform. This is the situation today; EMI and IGE represent Technology Providers of the predominant Grid middleware software deployed in EGI’s production infrastructure. It is important to note that EGI’s primary service to the Research Communities is the deployment and operation of their respective Community Platforms on their behalf. Typically, this scenario is called “managed hosting” in the commercial sector.

Since Research Communities own their respective Community Platform, it is their primary responsibility to *also* enter a partnership with Technology Providers for the Software maintenance and development activities.

Through the second pillar of its strategy, EGI provides community and coordination services (not shown in this figure) for Research Communities *and* Technology Providers. Frequently combined with technical services coming from the EGI Collaboration Platform, these services include Requirements management, Release Management, Software Vulnerabilities Management, Help Desk and Software Support, strategic Platform evolution, IT Service Management, etc.

¹ These services are often support services referred to as “EGI Global Tasks”.

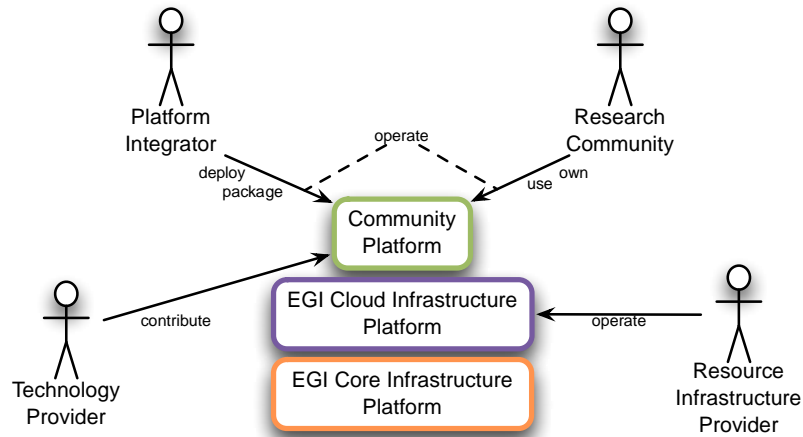


Figure 4: Relationships around Community Platforms deployed in EGI's Cloud Infrastructure Platform

The situation changes when Research Communities start integrating their Community Platforms with the EGI Cloud Infrastructure Platform, as shown in Figure 4. The biggest difference compared to the situation described before is that EGI's managed hosting service is no longer a required part of the service offering: It is typically the Research Community itself who operates their Platform through the consistent management interface exposed by the EGI Cloud Infrastructure.

Additionally, Figure 4 illustrates a variation of the basic scenario where a Platform Integrator takes care of the technical integration of software components with each other, and in this case the EGI Cloud Infrastructure platform. It is expected that some Research Communities will seek combined outsourcing of Community Platform integration and operation on top of the EGI platforms. Discussions with Technology Providers at the EGI Community Forum 2013 in Manchester indicate that there is interest among these stakeholders for this business model. This scenario is the blueprint for a marketplace for a rich set of services surrounding Cloud Computing, where each stakeholder offers their services in the EGI ecosystem.

Again, for clarity, this diagram does not show a number of support services. Strategic technical coordination and collaboration is continued to be offered through participation in the EGI Technology Coordination Board (TCB), and technical release management and provisioning services are available through participation in the newly formed EGI UMD Release Team (URT). Necessary revisions to the TCB Terms of References are underway at the time of writing, and a first draft of the URT Terms of References is currently being discussed.

2.3 EGI Core Infrastructure Platform

The EGI Core Infrastructure Platform (CIP) consists of a number of components that are necessary to federate Distributed Computing and Data Infrastructures into a (set of) consistent resource access services across administrative domains (nationally or globally).

This platform is deployed directly on top of the physical hardware owned by the Resource Infrastructure Provider in the management infrastructure part of the provider's physical infrastructure. The EGI Core Infrastructure Platform is owned, deployed, packaged and operated by the Resource Infrastructure Providers federated into EGI. These activities are coordinated through management and collaboration boards and groups that are part of the EGI IT Service Management domain.

The EGI CIP comprises of the following seven components:

- Messaging
- Federated Authentication and Authorisation Infrastructure (AAI)
- (Service Availability) Monitoring
- Accounting
- Central Services Catalogue
- Information Discovery Service
- Metrics visualisation

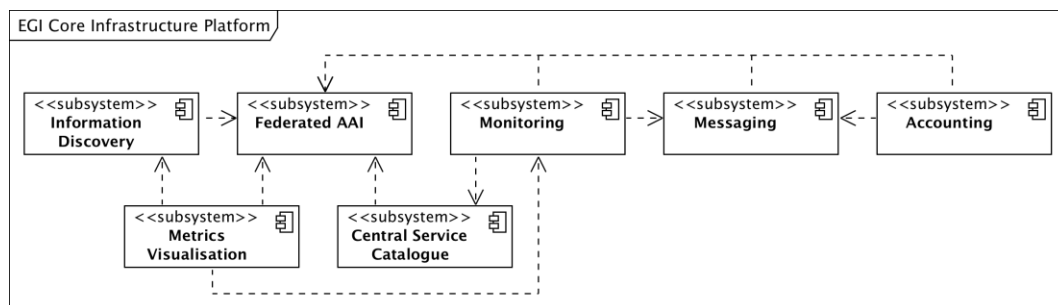


Figure 5: Components of the EGI Core Infrastructure Platform

Figure 5 illustrates the principal composition of the EGI Core Infrastructure Platform. The composition of the diagram indicates components that form the integration surface of the EGI CIP, and which are designed to be operational services (i.e. Metrics visualisation and Central Service Catalogue). At the time of writing not all five available integration components (Information Discovery, Monitoring, Federated AAI, Messaging, Accounting) are compulsory for Community Platforms to integrate with. Currently, Community Platforms are expected to integrate with the Federated AAI, Monitoring and Accounting components, while the Central Service Catalogue will have to accommodate for the registration of the Community Services. From a technical perspective, the Messaging component is used only as a sub-component of Accounting and Monitoring, but it has in and by itself potential as a first class component in the EGI Core Infrastructure Platform, providing value to *any* Community Platform deployed on top of EGI Platforms. A specific service model would need to be developed, but is expected to align with provisioning specific messaging topics and queues according to agreed service levels (e.g. throughput, retention, reliability, points of presence).

2.4 EGI Cloud Infrastructure Platform

At the heart of the EGI Cloud Infrastructure Platform (CLIP) are locally deployed Cloud Management stacks. In compliance with the Cloud computing model, the EGI CLIP does not mandate deploying any particular or specific Cloud Management stack; it is the responsibility of the Resource Providers to research, identify and deploy the solution that fits best their individual needs for as long as the offered services implement the required interfaces and domain languages.

Consequently, the EGI Cloud Infrastructure Platform is built around the concept of an *abstract* Cloud Management stack subsystem that is integrated with components of the EGI Core Infrastructure Platform (see Figure 6).

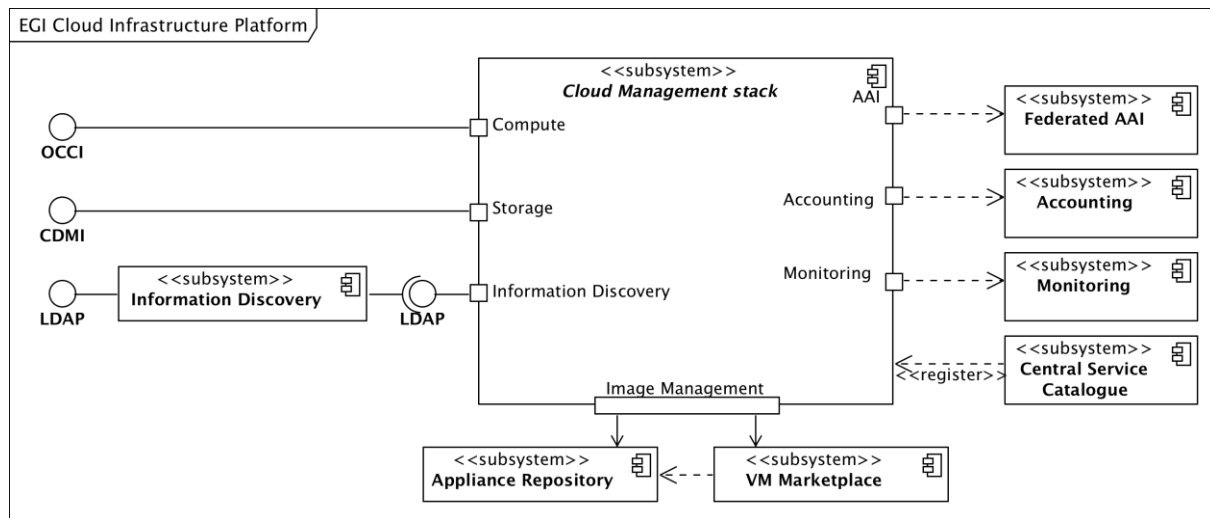


Figure 6: Architecture of the EGI Cloud Infrastructure Platform

As already indicated in section 2.2 the EGI CLIP is owned by the EGI Resource Infrastructure Providers through EGI. However, taking the nature of Cloud Computing into account, allows EGI to define the CLIP as a relatively thin layer of federation and interoperability around local deployments and integrations of Cloud Management stacks.

This architecture is modelled using an *abstract* Cloud Management stack subsystem (see Figure 6) defining interaction ports with a number of services from the EGI Core Infrastructure Platform, and the EGI Collaboration Platform. At the same time, it defines the required external interfaces and corresponding interaction ports. All these ports will have to be realised by local Cloud Management stack deployments.

The main interaction points of the EGI CLIP thus are:

- Integrate with the EGI Core AAI
- Integrate with the EGI Core Accounting system
- Integrate with the EGI Core Monitoring system
- Provide a standardised Cloud Computing interface (OCCI)
- Provide a standardised Cloud Storage interface (CDMI)
- Provide a standardised interface to an Information Service

Additionally, by means of using the Appliance Repository and the VM Marketplace from the EGI Collaboration Platform (COP) the EGI CLIP is providing VM image sharing and re-use across EGI Research Communities.

At the time of writing the EGI CLIP is maintaining its own, separate Information Discovery system. Even though it is using the GLUE2 schema, some extensions and tweaks are not compatible with the canonical GLUE 2 specification. Hence Cloud Resource providers maintain local LDAP endpoints (usually deployed as a resource BDII) aggregated into a Cloud Platform Information Discovery service, which in turn allows access to the data using LDAP v3.

Figure 7 provides an overview of the current realisations of the abstract Cloud Management stack subsystem in the EGI CLIP. It illustrates that each existing realisation inherits the obligation to implement the interaction points from the generalised parent Cloud Management stack. At the same time, the EGI Federated Clouds Task (funded through the EGI-InSPIRE project) gives Resource

Providers a platform to share their implementation solutions for a commonly deployed specific Cloud Management stack (e.g. OpenNebula and OpenStack).

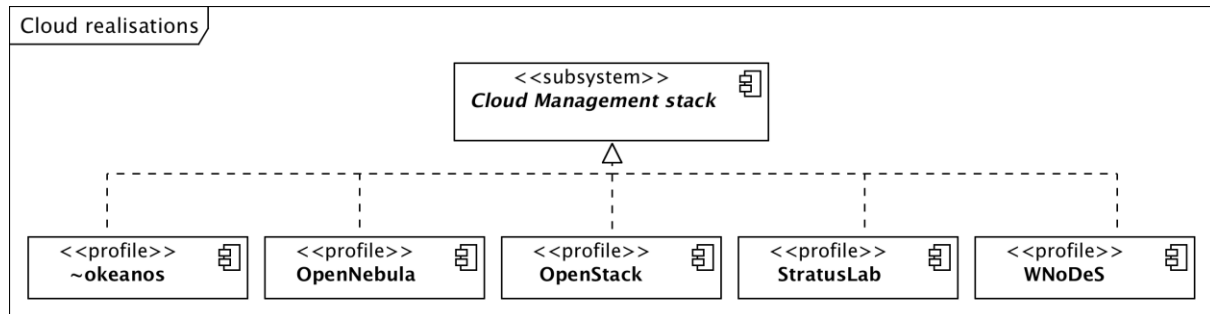


Figure 7: Current realisations of the abstract Cloud Management stack component

Through this collaboration, Resource Providers gradually develop and mature deployment and configuration profiles around common Cloud Management stacks as illustrated in Figure 7. These profiles will be documented and shared as blueprints available to the public during PY4.

2.5 EGI Collaboration Platform

The EGI Collaboration Platform (COP) is a collection of components that facilitate collaboration in the EGI community. Broadly, these services can be further structured into two categories of services:

- Social collaboration, and
- Technical collaboration.

While technical collaboration tools are designed and offered for integration with technical components of any number of Community Platforms deployed in EGI’s production infrastructure, the social collaboration components are typically tools and services that require direct user interaction (see Figure 8).

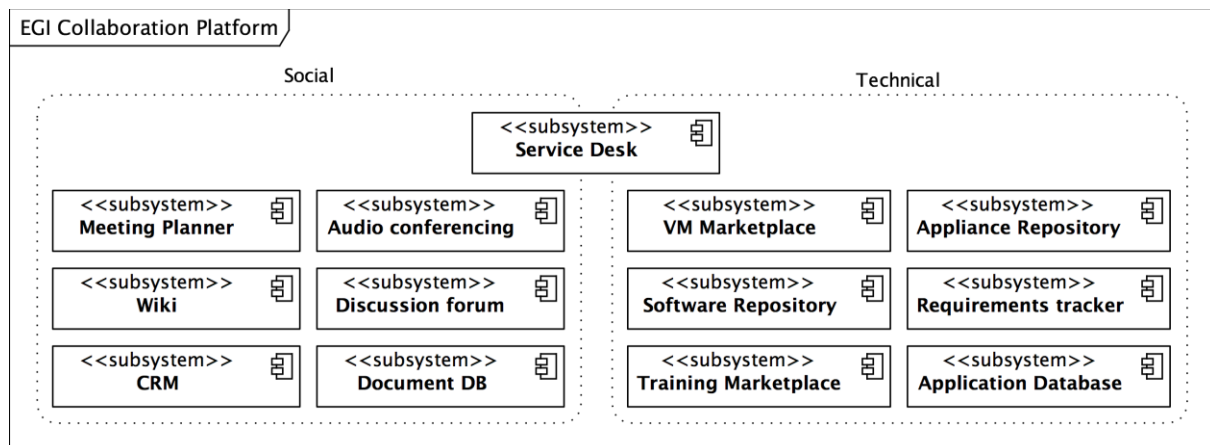


Figure 8: EGI Collaboration services capture social and technical collaboration needs



2.6 Community Platforms

The EGI Platform Architecture allows any number of Community Platforms being deployed in EGI's production infrastructure. Some may be deployed on top of the EGI Cloud Infrastructure Platform, while others may be deployed directly on top of the EGI Core Infrastructure Platform (i.e. a direct hardware deployment, yet integrated with the EGI Core Infrastructure Platform).

To date, EGI has traditionally deployed Grid middleware that focused on very few User Communities that were very well aligned in their requirements on a distributed computing infrastructure. The EMI collaboration integrated three major Grid middleware (gLite, Arc and UNICORE) and dCache into a regularly updated Grid distribution. This was then provisioned using documented and repeatable procedures in the EGI production infrastructure. The IGE software stack, which is essentially a Globus Toolkit augmented with additional services and integrations tailored for the European research communities, together with a number of alternative Grid middleware stacks now form the basis for EGI to expand further on formalising relationships with its Technology Providers.

Over the last year, EGI has started to develop not only the EGI Platform Architecture (embedded in this document) but also a support model for Technology Providers in the EGI ecosystem [R 8]. That document expands the definition of a Technology Provider in that it can act as a Platform Integrator, a Product Team or a very small, volunteer-based collaboration without agreed central coordination. While Platform Integrators are expected to provide complete Community Platforms, integrating a number of tools and services with each other (and within the platform) and with the target EGI platform (Cloud or Core Infrastructure), Product Teams typically provide one or more distinct components for integration into specific community platforms. Overlapping with Platform Integrators, Product Teams may also provide Community Platforms². Volunteer teams typically provide smaller components that are integral parts of other platforms.

With the end of the EMI and IGE, EGI needs to consider the continuation of its services and deployed platforms beyond these projects. What was deployed as one Community Platform, the UMD, before, might in the future separate into much more individual Community Platforms next to a probably much smaller UMD. Currently, EGI is considering engaging in closer conversations on future collaborations with a number of Technology Providers emerging from the technology coordination in EGI over the past three years, as indicated in Table 2.

Platform/Product	Lead institute	Key Capabilities	Included Products
ARC HTC	Lund University	HTC Compute, Info Discovery, Client tools	ARC CE, ARC Clients, ARC InfoSys, GridFTP wrappers
Globus	EGCF (c/o LRZ)	HTC Compute, Credential Mgmt, ...	GRAM5, Gridway, GridSAFE, Globus Default Security, GridFTP, MyProxy, GSISSH, GSISSHTerm, OGSA-DAI
UNICORE HPC	FZ Juelich	HPC Compute, File Transfer, Storage	UNICORE Gateway, UNICORE XUADB, UNICORE UVOS,

² These Product Team Community Platforms are typically very focused around a specific and closely related number of capabilities, e.g. Storage.

Platform/Product	Lead institute	Key Capabilities	Included Products
			UNICORE TSI6, UNICORE/X 6, UNICORE Client, UNICORE WS, UNICORE HILA
Data-intensive HTC	INFN	(Parallel) Compute, Client tools, WN profiles, Notification, Credential Management	CREAM, MPI integration, WN, UI, L&B, Proxyrenewal
QosCosGrid	PSNC	Cross-cluster computing, Multi-scale computing, Broker, Notification, Client Gateways	QCG-Compute, QCG-Notification, QCG-Broker, QCG-ScienceGateways
dCache	DESY	Storage	dCache
StoRM	INFN	Storage	StoRM
DPM	n.n.	Storage	--
Hydra	n.n.	Credential storage	--
AMGA	n.n.	Metadata Catalogue	--
WMS	INFN	Job Scheduling/Brokeri ng	--
FTS	CERN	File Transfer	--

Table 2: Provisional list of Technology Providers

3 EGI PLATFORM ROADMAP: SUPPORT MODELS AND PROCESSES

The EGI Platform Roadmap in its first iteration [R 3] has provoked discussions and as a result a more focused definition of both the technical and IT Service components, have taken place over the last year.

These have cumulated in a working document that focuses on the service offerings for Technology Providers [R 8], which has been already referenced earlier in this roadmap. It follows a pattern that is applicable to any service offering that is being developed and made available to consumers, as illustrated in Figure 9.

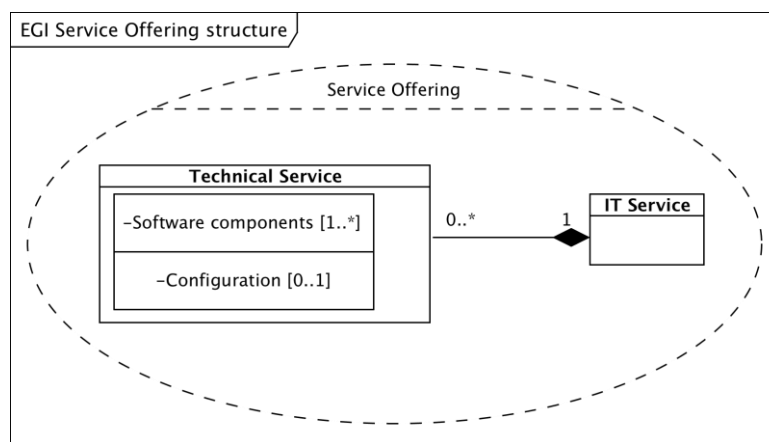


Figure 9: A typical structure of IT Service offerings

In this pattern, any IT Service is a composition of a number of technical services, which in turn comprises of a number of well-defined software components (typically software services). It is often the case that different IT Services reuse the same software components, but in different editions, or configurations, as illustrated in Figure 9. Software components that are designed to support different configuration profiles³ are particularly well fitting in such patterns and can be included with minimal effort.

Although rare, some IT services may entirely not require any backing technical service, and thus can stand on their own. These are often services in the community/social domain, providing dissemination, marketing, administration or coordination value.

Finally, any number of IT Services comprise together one of the principal service offerings provided by EGI.

Within this context a number of activities need to take place over the next year (which is at the same time PY4 of the EGI-InSPIRE project) for EGI to move towards a sustainable, generic, yet flexible e-Infrastructure service. The following subsections illustrate these activities in a greater detail. Nonetheless, these sections describe the activities on a generic and high level across all platforms. The actual realisation will be described in much greater detail in the EGI Technical Roadmap [R 4].

³ Configuration profiles in this context are not restricted to configuration files with different content per se. Often, configuration files are very complex in the attempt to provide maximum flexibility, frequently causing unnecessary configuration mistakes. On the other hand, different software editions may result in a different binary composition of the software product that cannot be captured otherwise in plain editable configuration files.



3.1 Maturing EGI Platforms

Defining the EGI Platform Architecture requires grouping existing and new components into different platforms. At the same time, stakeholders and roles need to be defined to provide more clarity on the roles and responsibilities of the participants in the EGI ecosystem.

The first edition of the EGI Platform Roadmap mainly focused on the early concepts of the EGI Platform architecture, a number of case studies and an attempt to define the Community Platforms deployed on top of the EGI Platforms.

EGI will need to define initial, yet stable versions of the EGI Platforms in order to allow building IT services and service offerings around them.

Once these start to mature, EGI will carefully look into the lifecycle management of platform components, using techniques such as technology insertion plans, technology assessment, pre-production procurement, Technology Provider procurement, etc.

3.2 Developing IT Services and Service Offerings

For each EGI Platform (i.e. Core Infrastructure, Cloud Infrastructure & Collaboration Platform), IT Services need to be developed that encapsulate a number of software services at given service levels into concise and coherent activities. Frequently, these IT Services may cross platform borders (for example, a Cloud Service may include VM sharing capabilities provided by the Collaboration Platform), introducing functional dependencies.

These need to be aligned with the needs of the customers of these services and platforms. While most of them will be designed around technical services, addressing Community Platform providers, some of them are addressing direct end user needs (e.g. an IT service developed around the Audio Conferencing system).

3.3 Engaging with Technology Providers

A key activity in the upcoming year will be engaging with Technology Providers for existing and new Community Platforms. This requires approaching User Communities, Software Engineers, and NGIs that are interested in engaging in this model of software architecture and support. A number of Technology Providers will emerge from the end of the EMI project as indicated earlier in this document (e.g. UNICORE, dCache), some exiting Technology Providers will continue to collaborate with EGI (for example, PSNC for the QosCosGrid platform) and some may collaborate with EGI under different conditions (e.g. the European Globus Community Forum, the successor of the IGE project).

3.4 Implementing new support structures

Part of the EGI Platform Architecture and Roadmap activities was an analysis of the current support activities and management boards, including conclusions on necessary changes to continue some of the support services that were provided by the EMI and IGE projects. The result of this was the decision to pursue a split of activities in the TCB, focusing the TCB more on strategic technology evolution, and combining the day-to-day coordination work in a newly formed coordination body named UMD Release Team (URT) which will replace some of the technical coordination functions that will be lost with the end of the EMI and IGE projects.

Draft Terms of References for the amended TCB [R 9] and new URT [R 10] were published and circulated among other EGI management boards (e.g. OMB, UCB and SVG) for comments. These



documents need ratification from the EGI.eu Executive Board, before the new management structure can be implemented.

3.5 Define IT services provided by Technology Providers

The further formalisation of EGI's relationship with Technology Providers will result in the definition of service targets for services that EGI needs from Technology Providers and, for those Technology Providers that wish to enter a formal agreement with EGI, in the definition and agreement on service levels for specific service targets.

Respecting the EGI ecosystem and nature of Technology Providers in that ecosystem, EGI will develop these in line and according to its service offerings for Technology Providers into offerings on three different engagement levels, provisionally named “community”, “contributing” and “integrated”.

3.6 Preparing EGI Platforms beyond EGI-InSPIRE

EGI has one year supported by EGI-InSPIRE to continue its transition towards a sustainable operational model. Part of this is an assessment of the individual components of the EGI Platforms, their popularity, uptake in EGI Communities, operational cost, maintenance costs and compensation EGI is receiving for operating these services. Particularly the maintenance costs are linked to the software provisioning services of the contributing Technology Provider.

For example, services in the EGI Core Infrastructure Platform, and the EGI Collaboration Platform are currently maintained using funds from the EGI-InSPIRE project (e.g. the Training Marketplace, Application Database, Accounting infrastructure, Monitoring). During this process, EGI needs to identify exactly which of these services it wishes to continue, extend, or discontinue, in line with EGI's business model and funding situation.

In a next step, EGI needs to procure maintenance services for these components from Technology Providers. Ideally, the IT Service models for the components in the EGI Platforms will be the same as for Community Platforms; hence coordination and collaboration should take place in the TCB and URT as well. EGI will have to seek for Technology Providers that are willing to engage with it on the “integrated” level for those components that EGI considers essential to its operational function (irrespective it being part of any of the EGI platforms).



4 CONCLUSION

The EGI Platform Architecture has been established within the EGI community as an approach to categorising and promoting the community's services and capabilities. Representing the first pillar of the EGI strategy, the EGI Platform Architecture clearly structures the components of the EGI production infrastructure along common goals and objectives. It allows for a clear separation of roles and concerns for all stakeholders in the EGI community, providing guidance both on how EGI will continue to evolve on the technical layer, and duties and responsibilities for partners in the EGI community.

Over the last year, exposing the EGI community to the EGI Platform Architecture, and equally important vice versa, has shown that EGI is heading the right way. Demonstrated through the EGI Cloud Infrastructure Platform, the chosen level of abstraction fits well with the EGI community. Although on a technical level, the chosen formality of publically defined standards (such as OCCI) fostered flexibility and the emergence of micro-ecosystems in the response of the participating Resource Infrastructure Providers around existing Cloud Middleware stacks. Through this, the EGI Cloud Infrastructure Platform is a living demonstration of the complementary nature of formality and flexibility, if the abstraction level is chosen right.

The process of finding a well fitting abstraction level in the technical architecture for EGI needs to be applied to the community and coordination services in terms of finding the right balance between formality (e.g. Service Level Agreements) and flexibility (e.g. fostering community relationships) within the EGI ecosystem. This will be the overarching theme of the months to come.

As a first visible action towards striking the right balance, the existing activities in the TCB will be split and augmented; the TCB will be adjusted to focus more on strategic evolution, whereas the newly formed UMD Release Team will take care of day-to-day activities (i.e. tactical evolution).

Further community and coordination services that are strongly linked with these activities will have to be either adapted (e.g. the TCB, Software Provisioning) or newly implemented (e.g. the URT) to provide services as part of the second pillar in the EGI strategy.

While the technical architecture of the EGI Platform model progresses well in its maturity, the service architecture aspect (the IT services from the second pillar) needs to catch up to settle into sustainable services until April 2014.

Consequently, the roadmap for the EGI Platform Architecture puts a focus on developing these services further, including building further relationships with new Technology Providers, and formalising the collaboration across contributions.

5 REFERENCES

R 1	D2.30 – EGI Strategic Plan, https://documents.egi.eu/document/1098
R 2	D2.31 – EGI Technical Roadmap, https://documents.egi.eu/document/1094
R 3	MS510 – EGI Platforms Roadmap, https://documents.egi.eu/document/970
R 4	D2.33 – EGI Technical Roadmap, to be published
R 5	EGI Technical Forum 2012, Prague, Czech Republic, 17-21 Sep 2013, http://indico.egi.eu/indico/conferenceDisplay.py?confId=1019
R 6	Evolving EGI Workshop, Amsterdam, The Netherlands, 29-30 Jan 2013, http://indico.egi.eu/indico/conferenceDisplay.py?confId=1252
R 7	EC Horizon 2020 strategy, http://ec.europa.eu/research/horizon2020
R 8	Post-EMI/IGE support for Technology Providers, https://documents.egi.eu/document/1499
R 9	TCB Terms of Reference, https://documents.egi.eu/document/109
R 10	URT Terms of References, https://documents.egi.eu/document/1618
R 11	ITIL 2011 Glossary, http://www.iti-officialsite.com/InternationalActivities/ITILGlossaries_2.aspx
R 12	EGI Technology Glossary, https://wiki.egi.eu/wiki/Technology_Glossary_Contents