EGI-InSPIRE

EGI Platform Roadmap

EU MILESTONE: MS510

Document identifier:	EGI-InSPIRE-MS510-v6.docx
Date:	**11/04/2012**
Activity:	**SA2**
Lead Partner:	**EGI.eu**
Document Status:	**FINAL**
Dissemination Level:	**PUBLIC**
Document Link:	https://documents.egi.eu/document/970

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 document describes the different platforms in more detail. The second part of the document provides a roadmap on adopting the platform-based architecture and service delivery model on the European Grid Infrastructure.

I. COPYRIGHT NOTICE

II. DELIVERY SLIP

	Name	Partner/Activity	Date
From	Michel Drescher	EGI.eu	5-42012
Reviewed by	**AMB& PMB**	EGI.eu	20-2-2012
Approved by	**AMB & PMB**		11-4-2012

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	5. 02. 2012	First draft	Michel Drescher, EGI.eu Alvaro Simon, CESGA
2	12. 02. 2012	Second draft for internal review	Michel Drescher, EGI.eu
3	19. 02. 2012	Third draft	Michel Drescher, EGI.eu
4	20. 02. 2012	Fourth draft, External review	Michel Drescher, EGI.eu
5	11. 03. 2012	Fifth draft, reflecting review comments	Michel Drescher, EGI.eu
6	18. 03. 2012	Sixth draft, reflecting moderator comments	Michel Drescher, EGI.eu
7	05. 04. 2012	Seventh draft, after AMB review	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 2] and the EGI Technology Clossary [R 3].

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 Infrastructure Platform	The EGI Infrastructure Platform comprises of IT Infrastructure and IT Services that are required by all Research Communities that are part of the EGI ecosystem in order to deliver community-specific services and infrastructure.
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.
EGI Community Platform	EGI Community Platforms (there may be more than one) consist of services that are specific to the respective community's needs.

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 illustrates how EGI may adopt a platform oriented IT architecture to deliver its services in a systematic way to a broad and diverse set of customers within the EGI ecosystem. By building on the current service interface, EGI will continue to deliver existing services to existing user communities while allowing its Resource Infrastructure Providers to broaden their customer base by offering new cloud related services integrated with some of the platforms described in this document.

The foundation of the EGI Platform architecture is the definition of the term platform itself. Defined as "[…]a composition of IT Infrastructure and IT Services that together enable independent solution providers to build other technologies or processes, or both, on top of it" (section 2) a platform's own added value will be its extreme horizontal scalability allowing many research community scoped value added services to build on top of it.

Thus, the EGI Platform model works with three distinct types of platforms that serve different purposes. The EGI Infrastructure Platform *enables consistent access* to a large federated distributed computing infrastructure comprising access to virtualised compute, storage and network resources, and supplemental services such as Information Discovery, Accounting, Monitoring and Notification that enables platform integrators to utilise this solid base to build different higher-level infrastructures (virtual research environments), for example targeting the requirements emerging from a wide variety of research infrastructures as documented by the ESFRI [R 1].

Extending the EGI Infrastructure Platform, the EGI Collaboration Platform *facilitates synergies* between Research Communities by providing services that are common across user communities without being domain-specific or critical to the operation of EGI itself. Services such as federated AAI (for Platform Users), Service Desks or meeting planning systems are good examples of such facilitating collaborative services.

On top of the EGI Infrastructure Platform, while making use of the EGI Collaboration Platform where required, any numbers of EGI Community Platforms provide *domain specific access* to the distributed EGI computing infrastructure and integrating elements of its platforms into domain specific Virtual Research Environments.

This document provides initial definitions of the EGI Infrastructure Platform and the EGI Collaboration Platform, and illustrates a set of viable EGI Community Platforms derived from EGI's existing infrastructure services. This will help the current EGI community start to focus on the composition of each platform, which community it will serve and how each platform will be supported in the years to come. It will provide the necessary impulse to EGI's virtuous cycle of continuous service improvement to extend to new ways of delivering an e-Infrastructure serving the future needs of the European Research Area [R 19].

TABLE OF CONTENTS

TABLE OF FIGURES

1 INTRODUCTION

EGI provides an e-infrastructure to support the data analysis and computational needs of its publicly funded and supported researchers across many diverse research communities within Europe and their partners world-wide. Allowing individual researchers and research collaborations to customise and therefore personalise the services they have access to when using EGI's resources is critical in broadening uptake across the diverse research communities that comprise the ERA.

Researchers need many ICT services to support the whole research lifecycle regardless of whether they work as individuals or in small or large research collaborations. However, the type of services that they require will vary depending on their research field and the scale of their collaborative activities. These services may range from the non-technical (e.g. bibliographic services, repository services, publishing services) to the technical (e.g. authentication services, data analysis services, workflow services, information services, data movement services) and the social (e.g. collaboration services, reputation services). These services need to scale either as individual instances or through interoperation with other instances across research communities of different sizes. EGI cannot expect to successfully scale its activities across all these areas. Therefore it must establish an ecosystem that allows the researcher (or those acting on their behalf) to provide a personalised e-Infrastructure for their use.

To satisfy these requirements this document outlines how a horizontal platform architecture helps achieving greater flexibility and efficiency in both provisioning and accessing distributed computing resources in a systematic way. However, this document does not in any form claim to provide a readily available solution for a sustainable future of EGI. It rather describes a *starting point* for discussions on the actual design and contents of future EGI platforms, and how and by whom they may be delivered.

Being the first of several iterations of the EGI Platform Roadmap, this document provides a definition of a platform as a combination of IT Infrastructure and IT Services, but continues to focus on the technical aspects of IT platforms delivered by the various stakeholders engaged with the different research communities.

The remainder of this document is organised in three parts.

Part one comprises of sections 2 – 4 and provides the foundation of the EGI platform approach. Section 2 clarifies the term platform and provides a definition that will form the basis for all subsequent sections, and for further iterations of this document. Section 3 introduces the stakeholders and actors that are interacting in this model. Section 4 completes part one with brief case studies illustrating various options available to the example Research Communities.

Part two describes in more detail the composition of the introduced platforms, with sections 5 focusing on the platforms owned and operated by EGI, and section 6 giving a brief overview of the EGI Community Platforms that may emerge out of the current Research Communities that make use of the EGI Infrastructure Platform.

Part three ventures into sketching how the near and medium term future may look like in section 7, and ends with a set of conclusions drawn from the other sections of this document.

2 DEFINITION OF A PLATFORM

Many different definitions of the term platform exist[1] and are often tied to a specific application area for which the definition is given. The most accurate yet generic definition of the term platform in an IT or Computer Science sense is probably this:

> *"A platform is any base of technologies on which other technologies or processes are built."* [R 4]

This definition lacks one important element that is often seen in real world platform deployment in that a platform is effectively the combination of technology and processes in a horizontal architecture that *allows* other *independent* technologies or processes to be built:

> *A platform is any base of technologies and/or processes on which other technologies or processes are built.*

Popular examples of that practical definition can be found in a large spectrum, from large commercial IT providers (e.g. Oracle Technology Network [R 5] or Microsoft Developer Network [R 6]) to non-commercial e-Learning platforms such as the Khan Academy [R 7] or the MIT OpenCourseWare platform [R 8]. They all share the commonality of composing technology platforms (i.e. IT infrastructure [R 2] and service platforms (i.e. IT Services [R 2]) to a single horizontal platform offering. The target domains for the various platform offerings, however, determine the actual mix and prevalence of offered services and technology.

That said, the term Platform in the EGI ecosystem is defined as follows:

> *In EGI, a platform is defined as a composition of IT Infrastructure and IT Services that together enable independent solution providers to build other technologies or processes, or both, on top of it leading to an added value for end-users.*

2.1 *Platforms and the EGI ecosystem*

In the EGI ecosystem, many actors and stakeholders collaborate and interoperate with each other on many different levels. In order to provide reliable, available, scalable and efficient access to resources and services across EGI, it is necessary to coordinate and deploy these resources and services in an organised manner. By organising these in distinct platforms EGI enables a much greater flexibility and independence of the various different platforms, thus may experience a significantly lowered coordination and integration effort when compared to the existing vertical service delivery model.

Therefore, platforms are scoped satisfying major concerns of the relevant stakeholder, and defined according to fundamental requirements that can be found on any level of abstraction across the whole software stack available in the EGI ecosystem. Well-designed platforms thus align well with the stakeholder's business models and fall "naturally" in place without considerable integration effort. The following subsections outline the technical aspects (i.e. the IT Infrastructure part) as an initial composition of platforms scoped around current and, insofar known, future needs of EGI's supported research communities. The IT Services elements of those platforms will be described elsewhere (e.g. future revisions of the *EGI Global Tasks review* (MS115 [R 22] or *Evolving the EGI Business Model*

[1] http://www.thefreedictionary.com/platform

[R 23]). Subsequent iterations of this document will continue to refine the definition of the various platforms in search for spot-on differentiations, adhering to the principle of *separation of concerns* [R 9].

A fundamental design aspect of the EGI Platform model respects the concurrent deployment (see Figure **1**) of the existing stacks of middleware services next to emerging development, deployment and operation of community platforms on top of the physical hardware that is managed by national Resource Providers federated into EGI (through NGIs).

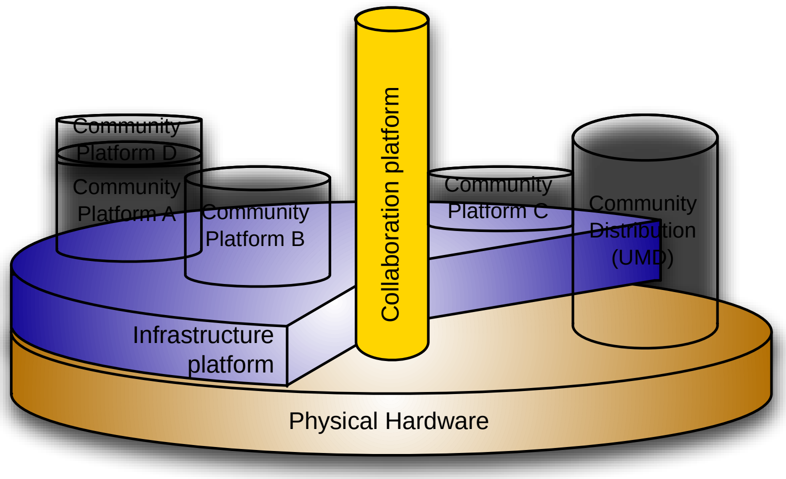

Figure 1: EGI platforms integrated with the EGI production infrastructure

2.2 *EGI Infrastructure Platform*

The EGI Infrastructure Platform's main scope is to enable flexible and efficient provisioning of IT resources irrespective of the customer's actual use of those resources. The EGI Infrastructure Platform will form the foundation layer of all other platforms that are, or will be, built on top of it. The IT Infrastructure and Services that will be part of the EGI Infrastructure Platform are those services that are required to successfully build a Community Platform on top of it. This is a logical evolution of the federated batch queue environment that has been brought into production over the last decade by EGI and its predecessors, where controlled remote access to private institutional batch computing resources for research communities is now being supplemented by access to private institutional cloud resources. It thus builds on top of the already federated computing resources within EGI and coordinated by EGI.eu.

The EGI Infrastructure Platform will be delivered as a federated service (IaaS) to its customers. Built on top of the existing physical hardware federated within EGI it exposes these compute, storage and networking resources as *virtualised resources*. This core service will be supplemented by further IT infrastructure helping platform integrators and operators to successfully build, integrate and operate appliances on top of the EGI Infrastructure Platform. The EGI Infrastructure Platform is also required to integrate with the existing set of services that are already deployed to manage and operate the current EGI production infrastructure.

The EGI Infrastructure Platform will be a core asset of EGI. Maintenance of the IT Infrastructure components of this platform should ideally be directly funded by the EGI community reflecting its need of a stable, manageable and efficient management infrastructure that persists across any funding project thus imposing a high-level (if not exclusive) control of the maintenance direction of the respective components. This will contribute to establishing a high degree of trust in the infrastructure by decoupling the funding stream from any time limitations imposed by a project-based cost recovery model.

2.3 EGI Collaboration Platform

The primary purpose of the collaboration platform is enabling the collaboration between communities that are using technology deployed on top of the EGI Infrastructure Platform. It will also build on top of the existing EGI production infrastructure, so that EGI's research communities are able to transition from the existing production infrastructure to integrating with the new EGI Infrastructure Platform, should they wish to do so.

The EGI Collaboration Platform comprises services and technology that are (expected to be) used across many if not all EGI research communities irrespective their scientific domain. The EGI Collaboration Platform therefore is supplemental to the EGI Infrastructure Platform even though all EGI Research Communities may use the offered services. The distinction between the EGI Collaboration Platform and any conceivable EGI Community Platform lays in the assessment of the technology relevance to the community's core IT business. Generic services (such as meeting management services etc.) may be popular among Research Communities, yet they are not part of their core infrastructure needs. Therefore it makes sense to include such services in the EGI Collaboration Platform. Specific services, particularly scientific applications (e.g. openModeller, used in the BioVeL project) clearly should be part of a Community Platform (e.g. for the LifeWatch Research Community).

Whatever the actual composition of the EGI Collaboration Platform, it will be an important EGI asset. Although collaboration services may play an important role, they are not considered part of the core IT Infrastructure needs (these are captured in the EGI Infrastructure Platform), Therefore, it will be delivered through a mix of direct funding, partnerships, sourcing in external SaaS offerings, or other collaborative activities.

2.4 EGI Community Platforms

EGI Community Platforms are best described as meeting the needs of the respective community. As a consequence, it is difficult to describe EGI Community Platforms in a generic way similar to the EGI Infrastructure Platform, or the EGI Collaboration Platform.

There may be considerable overlap in deployed services and applications between the EGI Community Platforms and the EGI Collaboration Platform. While one service may be offered as part of the EGI Collaboration platform (for example VOMS, or VM Image sharing) it is perfectly possible that the same service is also present in a community platform but in a different deployment or implementation, for a number of reasons. This is not considered as a problem. The research community itself defines the scope of their community platform, and therefore subject to the community's choice of software products to deliver the included services. While it may be obvious in such a situation to engage in synergy exploitation discussions, the EGI Platform model ensures that the involvement and impact can be kept independent from the maintenance and operation of the EGI Infrastructure Platform.

Community Platforms play a pivotal role in the EGI community, as illustrated in Figure 2. Whatever the actual deployment of any particular platform, the end-user experience should not change (as indicated by the identical height of the various different platforms). Traditional platforms are directly deployed on the physical hardware, just as the Grid Middleware service collated in the Unified

Middleware Distribution (see also Figure **1**) are deployed in the current EGI production infrastructure. Using current infrastructure management services such as Accounting and Monitoring (depicted in red in Figure 2) the services are delivered to expectations. The same services are part of the EGI Infrastructure platform, which serves Community Platforms that are stacked on top of it, i.e. fully exploiting the virtualised resources exposed by the EGI Infrastructure Platform.

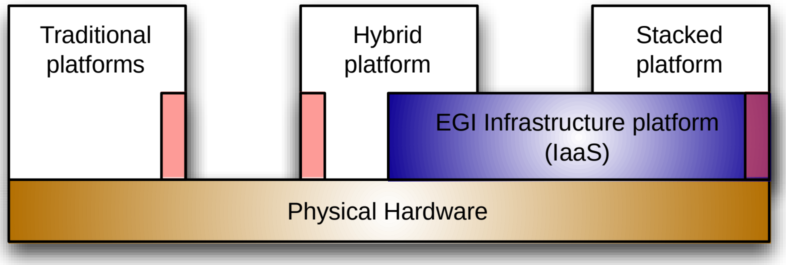

Figure 2: Different deployment models sharing operative management (in red)

As a third model, hybrid platform deployments make use of direct deployment on physical hardware. For example, a Cloud Provider may decide to offer supplemental services, as a means to distinguish itself from other members of the EGI Infrastructure Platform, to a selected set of Research Communities, and therefore decides to host a set of community-specific services (that nonetheless are part of the respective community's Platform architecture).

Unlike the EGI Infrastructure Platform and the EGI Collaboration Platform, EGI Community Platforms are expected to be an asset of the respective research community. The initial assembly and further maintenance of EGI Community Platforms should be directly funded by the owning research community, or shared between EGI communities, perhaps through forming EGI Community Platform consortia, forums or any other means of facilitating collaboration. The actual effort of assembling and maintaining an EGI Community Platform does not necessarily have to be adduced by the owning EGI Community itself. The community's steering body may decide to fund external effort (e.g. by contract, project or collaboration) to develop, package, and integrate its own platform as it may see fit.

3 STAKEHODERS AND ACTORS

Transitioning from a vertical service delivery model to a horizontal platform deployment model requires re-examining and identifying stakeholders and actors that together deliver end-to-end services for the research communities. This ensures that also requirements of those are addressed and satisfied in the newly proposed EGI ecosystem.

While it was possible to "hide" the distinction of roles and stakeholders in a vertical service delivery model, service platforms enable a much greater independence of actors operating at the various platforms built on top of each other. By carefully defining and scoping the different roles and identifying relevant stakeholders in a horizontally organised EGI ecosystem identification of business opportunities and orchestration of the various activities become much clearer. Synergies can be thus much more clearly leveraged.

This does not mean, however, that all discussed stakeholders and actors will have to remain separate entities as illustrated in

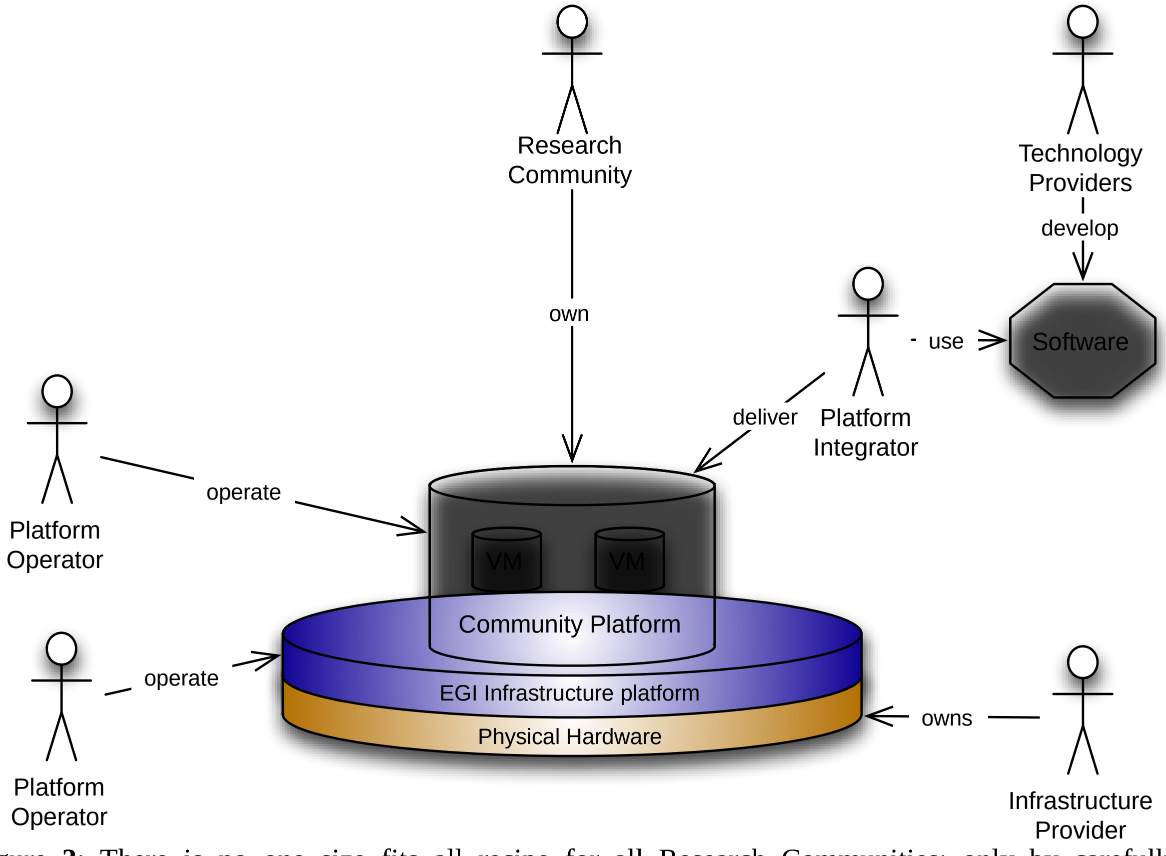

Figure **3**: There is no one size fits all recipe for all Research Communities; only by carefully examining the business needs and options (based on information given in D2.18: Evolving the EGI Business model [R 23]) any Research Community will be able to identify which of the roles (stakeholders and actors) it decides to assume, and which will have to be filled out by others in the EGI ecosystem. While it is expected that the larger a Research Community is (or grows into over time), the more roles it tends to assume, this will have to be proven by reality in the future.

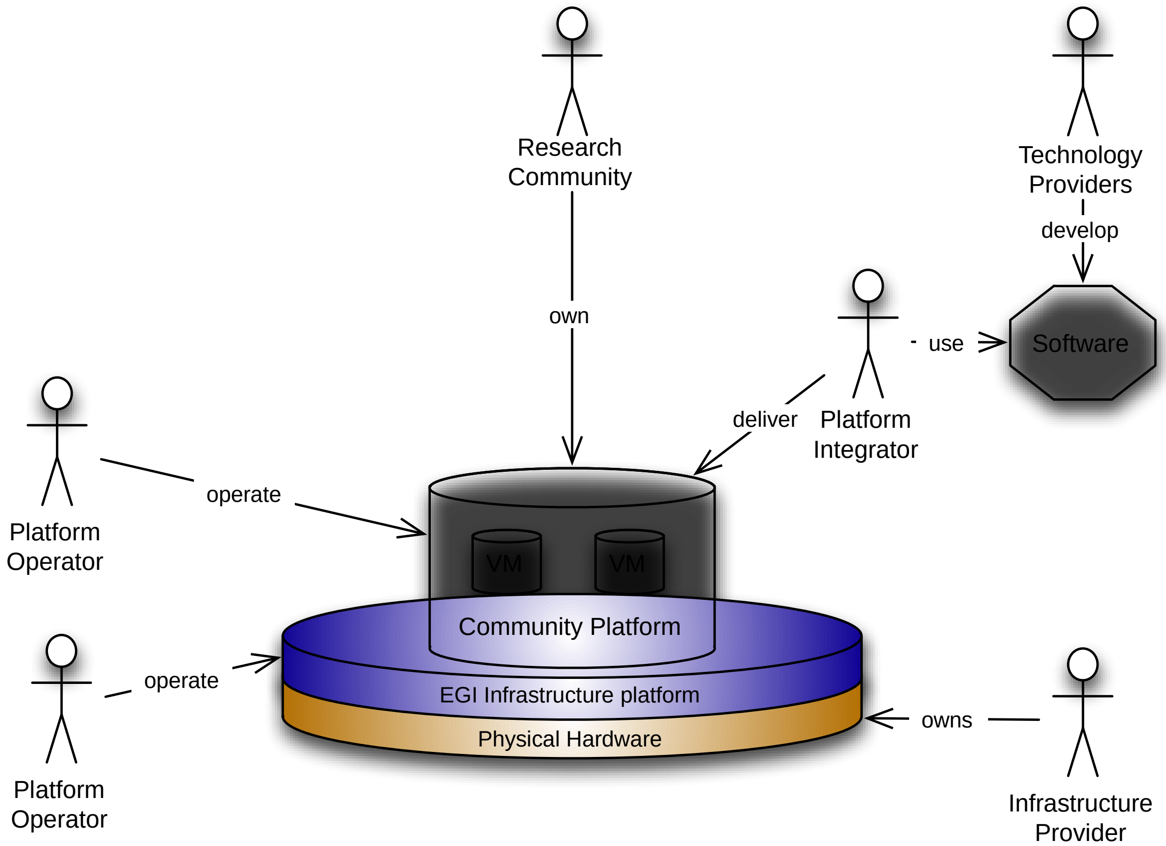

Figure 3: EGI community stakeholder interactions.

3.1 Stakeholders

3.1.1 Technology Providers

Technology Providers develop software according to available requirements and needs. Architecture and design of the software may be community specific, or of general purpose that may serve any consumer. Depending on their involvement in the EGI ecosystem, business and service models, or the software's main aims, the actual interest in the EGI ecosystem may vary across Technology Providers.

For example the Apache Foundation may have very little interest in the EGI ecosystem as such, yet it must be considered as a Technology Provider in the Platform Integrator's choice of software suppliers. On the other hand a Technology Provider may have strong interests in providing software tailored to the needs of a supported community. The EGI Application Database [R 17] provides many examples of applications written and maintained by Technology Providers with dedicated, specific community scope. Current Technology Providers for the EGI community, such as EMI and IGE provide software that is deployed in the current EGI production infrastructure. In a concurrent deployment scenario the current scope and role for EMI and IGE as Technology Providers for EGI may persist, while in a scenario utilising the EGI Infrastructure Platform, the sustainability of a subset of currently provided services may be limited.

Typically, Technology providers deliver software as source code, binaries compiled for a specific execution platform, or both (just like the software registered in EGI AppDB today). Delivery of software ranges from online code repositories that are either self-managed or externally managed (e.g. Sourceforge[2], GitHub[3] or Launchpad[4]) to shipped media (CDs, DVDs, USB sticks, etc.). Support for the provided software varies greatly, depending on the Technology Provider's business model.

Engagement with Technology Providers may happen at all platforms present in the EGI ecosystem, from the EGI Infrastructure Platform to the various EGI Community Platforms built on top of it.

3.1.2 Platform Integrators

Platform Integrators architect and design a platform against identified requirements delivered by their customers. During that process Platform Integrators match the requirements against *available* software and select the most suitable software components according to additional criteria (such as ease of customisation, configuration). An important aspect of this platform design process is the actual selection of a suitable lower-level platform to integrate with. Depending on customer requirements, available software, engineering skill sets and licensing models (next to many other potential selection criteria) a Platform Integrator may choose one, or many lower-level platforms for integration.

In an ideal world, a Platform Integrator may choose from available software that behaves perfectly well according to documented interfaces and deployment guidelines. In reality, however, this is often not the case, and software "glue" (e.g. adapters for certain incompatible functionality) is required to be able to integrate two components. That glue software is hidden, and not included in the official external public platform interface and documentation. The extent of the required integration effort has a strong influence on the selected components, ranging from near to zero integration effort of perfectly interoperable components to significant integration effort for components that are used beyond their original intent.

The IGE project currently fits the role of a Platform Integrator. Integrating and adapting the Globus Grid Middleware for specific needs and constraints in the ERA (e.g. privacy requirements) IGE provides building blocks for a PaaS on top of which Research Communities deploy their own specific applications.

3.1.3 Platform Operators

Platform Operators – as the name suggests – operate a deployed and initially configured platform on top of its selected infrastructure. This day-to-day activity includes monitoring the platform infrastructure, administering changes if operational metrics are outside of acceptable upper or lower bounds, and reporting for pro-active platform provisioning.

Platform Operators are therefore interested in platforms that are easy and efficient to manage. Their requirements on a platform focus on scalability, reliability, accuracy and efficiency of the platform management infrastructure.

3.1.4 Resource Infrastructure Providers

Infrastructure providers are those partners in EGI who provision the tangible resources (compute, storage and network) in the EGI Infrastructure Platform providing transient access to these resources for a certain amount of time to a known set of users.

---

[2] http://www.sourceforge.net
[3] https://github.com/
[4] https://launchpad.net/

Resource infrastructure providers carry the risks and responsibilities of ownership of those resources, but at the same time have the control on who they allow access to these resources. Resource infrastructure providers are interested in a sustainable customer base that does not threaten their business models should one or more customers terminate the business relationship. Therefore infrastructure providers require a platform that exposes their resources securely, yet allows for flexibility and uniformity irrespective of how customers are actually making use of the leased resources – tying one's business models into the customer's business models potentially threatens sustainability of an infrastructure provider as a whole.

3.1.5 Research Communities

In the EGI ecosystem, Research Communities are a group of collaborating researchers that sustain a distinct (perhaps dedicated) management function that coordinates activities within the Research Community, and maintains relationships with other, external stakeholders within or without the EGI ecosystem.

Research Communities pursue strategic goals for the benefit of the collaborating scientists and research projects the Research Community participates in. As such, Research Communities are interested in platforms that deliver exactly the functionality they need, and responds efficiently and timely to evolving needs.

3.2 Actors

3.2.1 Platform Packagers

Platform Packagers turn the documented architecture and design of a given platform into artefacts that can be deployed on a target platform. Depending on the scope and definition of the platform those artefacts may be binary code packages such as RPM archives, or a larger structure and set of packages that together deliver a service as part of the platform.

With that, Platform Packagers take care of the technical platform lifecycle. This begins with assembling and publishing the initial release of the platform as a whole, including the technical documentation. The packager then monitors the development activities within the individual lifecycles of the included components. If required the packager plans and initiates updates to the platform components, thus creating a lifecycle of their own for the deployable platform components.

Platform Packagers often re-use components in order to simplify the process of aggregating low-level functionality into a higher-level service. Re-using software component also reduces the number of dependencies and effort necessary to monitor and track the development of the selected components. On the other hand, Platform Packagers must keep an eye on the quality of the selected components in terms of software defects (software problems and vulnerabilities), since each re-use of a component raises the impact of any of those software defects in the deployed platform.

3.2.2 Platform Deployers

Often overlooked as a distinguishable role, the Platform Deployer takes care of rolling out the components of a chosen platform at a specific time and place and configuration. This is documented in detailed roll-out plans that align with generally planned maintenance cycles of the production infrastructure (or may warrant a specific maintenance cycle if required).

In due time the Platform Deployer then implements the planned roll-outs (or updates the deployment plans). Each rollout of a platform component is documented in the "roll-out history" of that component to reflect the most current configuration state of that element for post-rollout consultation and troubleshooting. This is often referred to as a "configuration item" whose current state is maintained in a configuration management database (CMDB) [R 10].

The role of a Platform Deployer is often assumed by Platform Operators since the topic of their duties is identical. However, viable scenarios separate those roles where the Infrastructure Provider conducts the platform deployment (as a service) while the consuming Research Community is assuming the role of the Platform Operator.

3.2.3 Platform User

Generally, all those individuals that access a given platform, or any type of software, are summarised as "end users". For the purpose of (at least) this document, these users are described as Platform Users. Platform Users (primarily from various research communities) use the chosen Community Platform; they consume its services and underlying resources, without maintaining the business relationships that make this possible.

In an "end-to-end" description of service delivery (see Figure **4**), Platform Users would be located at one end of the service value chain, while the Resource Infrastructure Providers are located at the other end of that chain. In a platform oriented service delivery model this notion does not change for the Platform Users, they still consume the services that were deployed for their direct use.

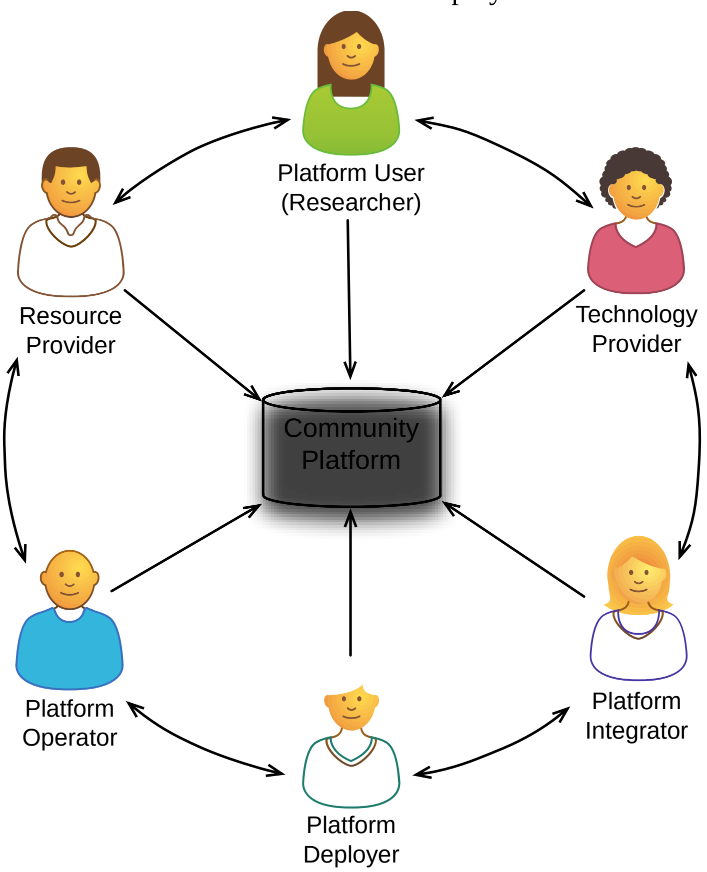

Figure 4: End-to-end value chain from Researcher to Resource Provider

This description does not imply a passive or receiving-only role. Instead, Platform Users are the main suppliers of functional requirements that reflect the needs of the respective community. By proxy, the Research Community ensures that the Community Platform, either by requesting a change to existing platform components or by having them replaced by a better alternative, meets these requirements. These requirements are the main drive the virtuous cycle of continuous service delivery to the Platform Users.

4 CASE STUDIES: EVOLVING EGI INTO A PLATFORM DRIVEN ECOSYSTEM

When studying the idea and potential of a platform oriented approach, the reader may wonder how the transition to such a model may look like. The following subsections describe how a transition to a platform based EGI could take place for selected EGI research communities.

The studied communities are by no means representative; neither are the transitions outlined in the respective sections. They simple represent some of the many possible options communities may have – communities are entirely free in selecting their own evolution path. The following scenarios remain on a high-level of description; they intend to provide a starting point for further, more detailed discussion within the various EGI communities.

4.1 Worldwide LHC Computing Grid (WLCG)

The WLCG collaboration is by far the largest Research Community that has been – and still is – supported by EGI. Through the shared history of both WLCG and EGI (including EMI) a vertical service delivery model has evolved through the EDG and EGEE series of projects over the past decade. As such, the evolving options for WLCG may serve as a blueprint for other research communities in the EGI ecosystem,

Essentially, the deployed services were developed by and for WLCG as a self-serving community that was its own Technology Provider, Infrastructure Provider, Platform Integrator and Operator (Figure 5). In fact the WLCG collaboration still includes all stakeholders and actors that are described in section 3. In that sense, the WLCG community can be seen as a self-sustaining entity.

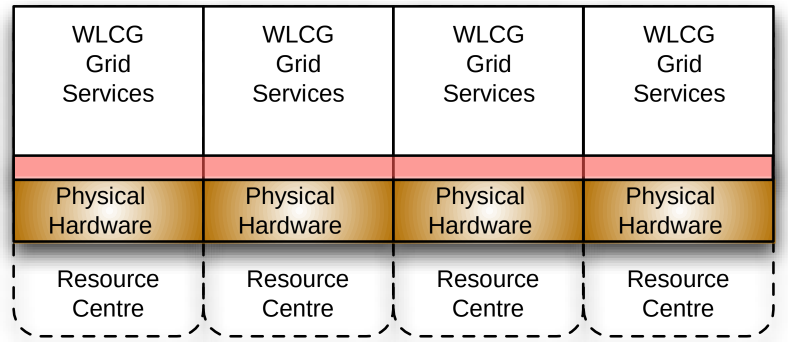

Figure 5: Schematic deployment of the WLCG Infrastructure.

The WLCG assets can be summarised at a high level as follows:

- WLCG – through the collaborating experiments - conducts a set of large experiments operating instruments of significant investment.
- Due to the nature of the experiments and the instruments, large amounts of data are produced and need safe, secure and efficient curation – WLCG supports a data deluge on their own.
- WLCG accumulated great expertise in data curation and distribution for world-wide access and analysis.
- The concept of Worker Nodes indicates an indirection in the architecture of the employed Grid middleware stack. By making all necessary Grid middleware client libraries available as

a single installation profile the actual local realisation of a Worker Node (e.g. by installing domain-specific applications) is abstracted away behind the Worker Node façade, allowing the Grid Middleware to develop independently from the applications interacting with them.

The following scenarios provide examples of how the WLCG may make use of the EGI Infrastructure Platform.

4.1.1 Use EGI Infrastructure platform as Worker Node infrastructure

In this scenario, domain-specific applications are encapsulated in Virtual Machines that are deployed and managed by the WLCG Grid Services – the Worker Nodes on Demand Service (WNoDeS) approach [R 11]. The Grid services remain being provided effectively as a PaaS service platform, supporting both traditional LRMS based cluster management and Virtualisation based compute infrastructure management through the WNoDeS abstraction layer.

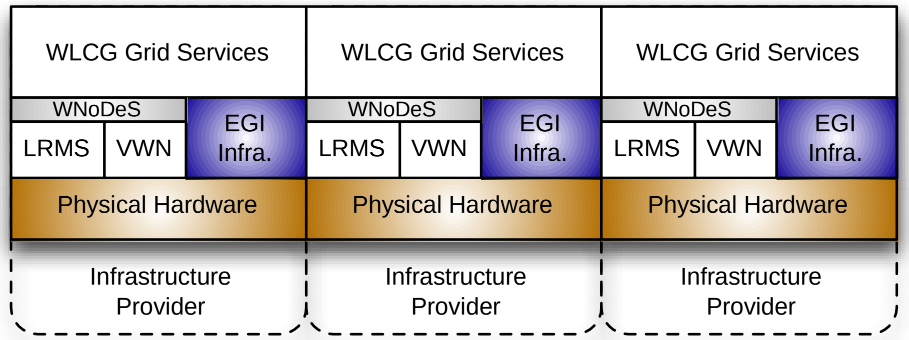

Figure 6: WLCG Grid services in a hybrid deployment.

Virtualising Worker Nodes may happen on two levels and in two phases allowing a controlled transition in the Grid Middleware: Traditional Worker Nodes are virtualised using WNoDeS technology, as a first step that stabilises VM image management, sharing and packaging processes transparent to the Grid middleware and overall WLCG platform operation. In a second step deployment and management of domain application extends to utilising the EGI Infrastructure Platform to deploy domain applications in an elastic infrastructure usage model, managed by Grid services (Figure **6**). Alternatively, both steps may happen simultaneously, allowing Grid services to transition independently from traditional LRMS based computing to IaaS based computing.

4.1.2 Outsourcing Software development and integration

This scenario illustrates the opportunity to formally and practically relinquish all software development and maintenance work to external providers. This scenario is not far from today's situation in that EMI is already an external Technology Provider for EGI serving the Heavy User Communities (HUC) in general. The difference lies in the setup and composition of the roles and responsibilities and the employed technology. WLCG effectively retains being a User Community (representing the LHC experiments), and a Platform Operator providing the WLCG Grid services to its associated research communities (see Figure **7**).

There is flexibility in this scenario regarding the role of Infrastructure Providers and delivery of software for deployment by Technology Providers. Two variants at the opposite ends of a scale of options may be used to gradually transition from direct deployment on physical hardware (current model) to full exploitation of Virtualisation as a deployment and operating technology (a possible future).

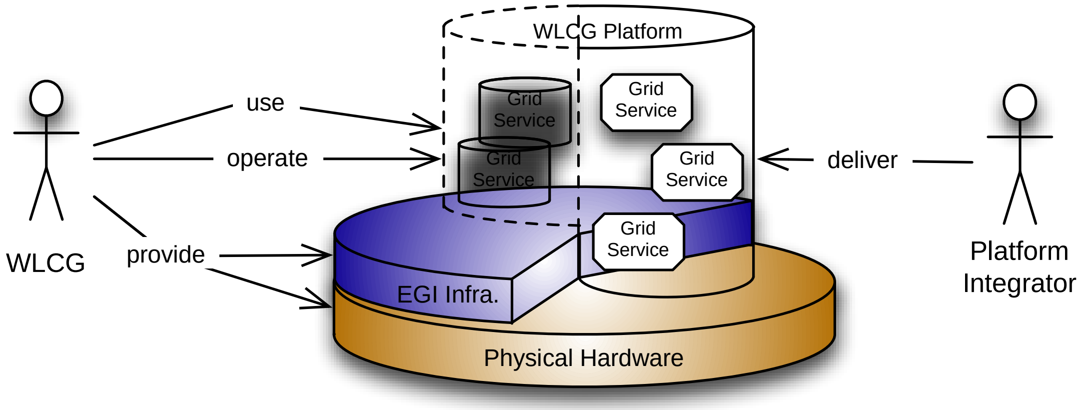

Figure 7: WLCG outsourcing Software development

4.1.3 Providing a Platform to its user communities

This scenario provides the opportunity for WLCG to focus on providing a tailored platform for its user communities while sourcing in the relevant resources from Technology Providers and Infrastructure providers. WLCG may identify and select a set of EGI Infrastructure providers that are able to satisfy its compute and storage demands according to negotiated SLAs, turning the underlying physical infrastructure into a deployment detail of the EGI Infrastructure Platform (Figure **8**). Likewise, Platform Integrators delivering necessary components to targets specified in another set of SLAs ensure that these components are packaged so that WLCG can deploy and operate the platform with as little configuration and contextualisation effort as necessary.

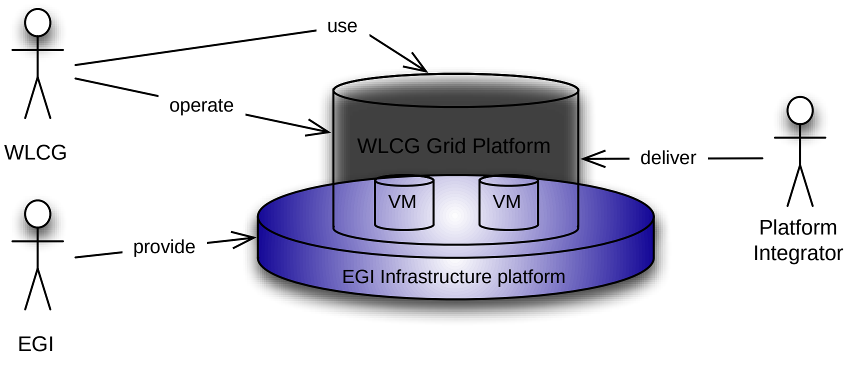

Figure 8: WLCG providing a coherent platform to the LHC experiments

4.1.4 Focussing on core business values and assets

This final scenario builds upon the biggest single asset of the WLCG community: Generating scientific value and progress through operating cutting-edge, community-tailored scientific instruments, i.e. the LHC (Figure **9**). This scenario allows WLCG to completely focus on their core strengths scientific data analysis, and to source in all other components from external service providers, such as EMI to provide the IT platform components, and EGI to deploy and operate the WLCG data analysis platform.

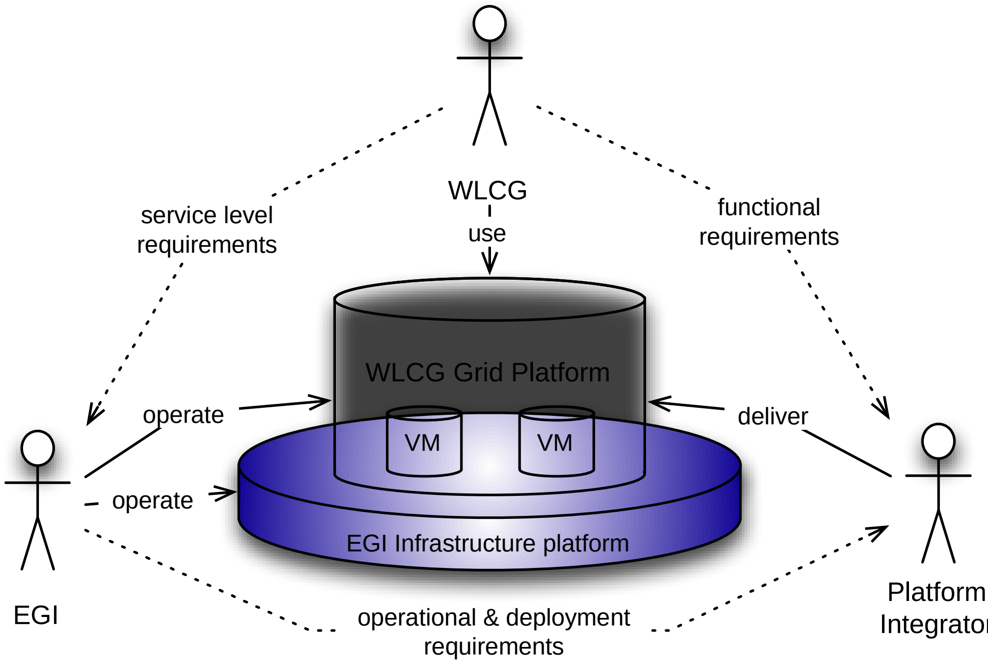

Figure 9: WLCG engaging in a network of services while focusing on its core strengths.

4.2 *BioVel*

The BioVel project explores how biodiversity research may efficiently utilise European e-Infrastructures in support of ESFRI projects such as LifeWatch. One activity examines and develops computational models for the biodiversity of species based on geospatial information (terrain topology, foliage distribution, climate, pollution) and factual recorded sightings, compute intensive simulations use this data as input to calculate species distribution probability maps using one of a pre-defined set of simulation algorithms.

The activity within the BioVel project is strongly focused on domain aspects of their research. IT infrastructure is mostly used as a commodity corresponding to the IT engineering capabilities available in the community. The community maintains the openModeller software hosted at SourceForce [R 12]. To access wide-scale distributed computing infrastructure, the community is pursuing a Virtual Machine based application deployment and operation model, reflecting the straightforward computing requirements.

However, the most apparent risk is the availability of IT engineering and operation skills within the community. It would be too costly for the community to develop these skills internally. Instead IT engineering and operation skills should be sourced from external, suitable sources.

In the EGI Platform model, the openModeller activity would mainly act as the Platform User of externally delivered services. Management functions of the associated Biodiversity Research Community coordinate business relationships to various providers of services that are not delivered through the BioVel project, or the Research Community itself.

Strategic alliances might be formed

- With experienced Platform Integrators, who may assemble a distinct set of VMs for the community ready for deployment and operation on the EGI Infrastructure Platform. Through the BioVel project, the community will continue to maintain openModeller by contributing to the application's project hosted at SourceForge.
- With Infrastructure Providers federated through EGI for accessing virtualised resources. Infrastructure Providers may also both deploy and operate the community's own platform on behalf of and as tasked by the community.

5 EGI PLATFORMS

Based on the definition and scoping of the EGI Infrastructure and Collaboration Platforms given in section 2, this section highlights the capabilities and services that may be delivered through the EGI Infrastructure Platform, and the EGI Collaboration Platform to provide the EGI communities with:

- Consistent, systematic and flexible access to virtualised computing resources, and
- Collaborative tools that help Research Communities to leverage potential synergies of using the same e-Infrastructure delivered by EGI.

A key property of the EGI Platform architecture is the federated delivery to the EGI ecosystem; there will be no single Infrastructure Provider, but a multitude of them federated together to ensure consistent access to their offered virtualised resources regardless of locality or consumed services.

It is important to differentiate Platform Operators from Platform Users. Although at times identical, a Platform Operator directly accesses the management services exposed by the EGI Infrastructure Platform, whereas the Platform User use the services exposed by the respective Community Platform, and only indirectly consume the EGI Infrastructure Platform services. This fundamental separation in the EGI Platform architecture allows a separated analysis and design of platform components as described in subsequent sections.

Software being part of a Community Platform that directly accesses the exposed EGI Infrastructure Platform services is included in this definition: From an EGI Infrastructure Platform point of view, it does not make any difference whether a human or a software service is accessing its offered services – both are operating the Community Platform through accessing the management interfaces exposed by the EGI Infrastructure Platform.

Further, the EGI Platform architecture reflects the expectation that the number of individual consumers of the EGI Infrastructure Platform is significantly smaller than the number of Platform Users, both across all deployed Community Platforms, and for individual Community Platforms. The general assumption is that Community Platforms are designed for sustainability, i.e. having a longer lifetime than the research projects that consume them. Therefore the fluctuation of Platform Users is expected to be higher that the fluctuation of Community Platforms, thus of Platform Operators as well.

As indicated in section 2.2, the EGI Infrastructure Platform is considered a core assed of EGI. Therefore, the Federated Clouds Task Force [R 13] is investigating its architecture and implementation of this platform as part of its remit. Consequently, the description of the EGI Infrastructure Platform is much more detailed compared to the EGI Collaboration Platform, and even more so when comparing it with the Community Platforms indicated n section 6.

5.1 EGI Infrastructure Platform

The EGI Infrastructure Platform (Figure **10Error! Reference source not found.**) enables flexible and efficient provisioning of IT resources irrespective of the consumer's actual use of those resources through adopting the Cloud Computing Infrastructure as a Service (IaaS) model delivered by federated Infrastructure Providers.

Virtualised resources (compute, storage and network) are delivered by the federated Resource Providers. This is supplemented by a set of services and technologies that together satisfy a wide variety of e-Infrastructure needs of European Research Communities. Each Cloud Management solution provides a Web UI that covers a visually accessible management functions for the virtualised resources. These are deployed locally and may differ significantly from each other. In contrast, the exposed APIs are required to expose the same standardised interface ensuring on the technical level consistent and federated access to the virtualised resources.

The following subsections describe the components of the EGI Infrastructure Platform in more detail.

5.1.1 Federated Authentication and Authorisation Infrastructure (AAI)

An integral element of the EGI Infrastructure Platform is to grant Platform Operators access to the infrastructure management services as described earlier. Depending on actual Platform Operator needs and Infrastructure Provider capabilities, the level of access to the exposed services will have to go beyond starting and stopping Virtual Machines. Therefore Platform Operators require strong authentication processes and security tokens to satisfy Infrastructure Provider needs. This aligns well with the expected long-term relationship between Infrastructure Provider and Platform Operator [R 20], where little fluctuation allows for more elaborate identity vetting processes.

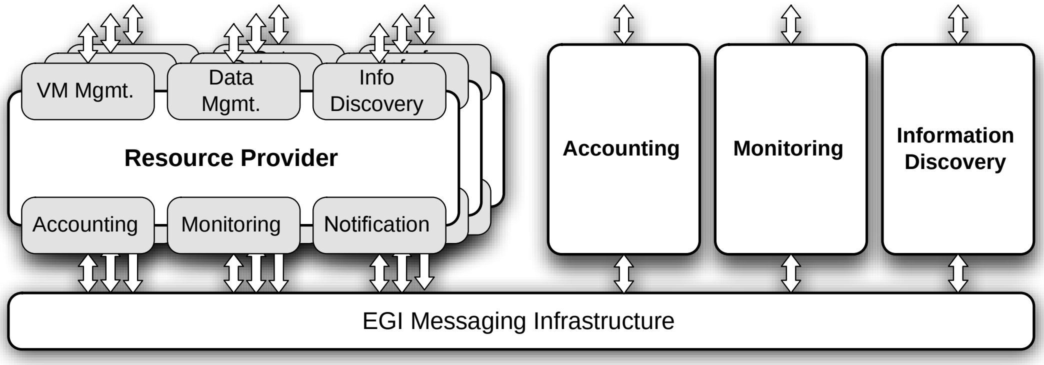

Figure 10: The EGI Infrastructure Platform architecture

X.509 based federated Authentication:

Platform Operators will be authenticated using X.509v3 certificates, re-using their existing Grid certificates. EGI already utilises a federation of Certificate Authorities by regularly adopting the set of trust anchors published by the EUGridPMA, and deploying them into the existing production infrastructure. Other solutions exist (such as eduGAIN, Shibboleth, InCommon), that may replace certificates in a federated Authentication infrastructure. However, the differences in the identity vetting process, and the potential of the named alternatives to hide a person's identity behind an opaque identifier may raise objections to their deployment in the medium term for allowing Platform Operators access to the EGI Infrastructure Platform management services.

Flexible role-based attribute authorities:

The platform must provide for a flexible and lightweight mechanism to declare a Platform Operator's membership or affiliation with a Research Community. These affiliation may be role-based, or of organisational nature such as VO membership, or both (e.g. acting as a Platform Operator for a given VO). The service that is maintaining these mappings must be accessible to the consumers of the EGI Infrastructure Platform for self-serving (i.e. declaring or ceasing a group membership), yet it must issue secure assertions when responding to an attribute request for a given identity.

Separating Platform Operators from Platform Users allows focusing on managing direct access to the EGI Infrastructure Platform services, unlike the current production infrastructure where access to Grid

Services (Grid middleware) was managed using proxy certificates to scale the service to approximately 18.000 researchers directly consuming the Grid infrastructure. SAML is a publically available standard that satisfies the requirements outlined in this section, and many different implementations exist.

A proven solution is available with VOMS, which is capable of issuing both SAML assertions and RFC proxy assertions upon request. This makes VOMS an ideal service enabling the transition from today's heavy utilisation of Proxy Certificates to a SAML based assertion infrastructure. While both solutions would satisfy the requirements illustrated in this section, only one solution should prevail in the long-term to reduce integration and maintenance effort for the EGI Infrastructure Platform. Also, considering the integration effort with other solutions (e.g. XACML, see below) for related capabilities and services, SAML presents itself as the better alternative for long-term sustainability.

Distributed policy-based Authorisation:

Complementing SAML as enabling technology for a flexible yet distributed authorisation infrastructure, XACML is the second enabling technology considered for the EGI Infrastructure Platform. As a combination, SAML and XACML allow a federated infrastructure to implement hierarchies of policies:

- At a global scope, e.g. implementing global banning policies,
- At a regional scope allowing NGIs to support select user communities with special conditions and access rights, and
- At a local scope enabling individual Cloud Providers take on User Communities on their own expenses and benefits.

A tool that is currently developed under the auspices of EMI is ARGUS, covering all-important aspects of a federated, distributed authorisation infrastructure. Being one of the younger solutions in the EMI product portfolio, its maturity is documented by other EMI products such as CREAM transitioning from its traditional authorisation means to integrating with ARGUS.

Selecting an enabling technology such as XACML intentionally introduces a dependency on a given technology on a global scale that must be well thought through, since all other services included in the Infrastructure Platform in turn are required to integrate with that technology.

XACML, however, is a mature standard maintained and evolved through OASIS and is supported by a great variety of both commercial and open source solutions, of which ARGUS is one example. Nonetheless, such a decision still requires EGI to assure itself of the sustainability of the ARGUS development in the future before the decision is finalised.

5.1.2 Cloud Management services

The core purpose of the EGI Infrastructure Platform will be providing access to virtualised resources. The three pillars of virtualised resources are *Compute*, *Storage*, and *Networking*. The goal is to provide consumers with a self-service management interface to configure all three virtualised resources according to the needs of the respective Research Communities.

Currently, mature services and interfaces are available for Computing, and to a lesser extent for Storage. No mature solution for self-service access to virtualised networking is available.

Possible solutions are OpenNebula, OpenStack or Eucalyptus representing the Open Source community. Commercial solutions are available, for example, from VMWare. It is important, however, that any considered solution supports public standards for Cloud computing. The EGI Federated Cloud Task Force [R 13] has identified OCCI 1.1 and OVF 1.1.0 as compulsory management interfaces for VM management, and CDMI 1.0.1 as candidate interface for management for virtualised storage.

5.1.3 Messaging

In a distributed environment that is used by many different components and services it is often very difficult to reach consensus on integrating on one particular protocol, or even a communication style (synchronous, asynchronous) to use for communication between services.

Messaging is a capability that encapsulates all details of communication between multiple participating endpoints, while providing important features that are required by many different business use cases such as delivery reliability, fail-over, synchronous and asynchronous message delivery, and asynchronous subscription management. This makes messaging a powerful component in a widely distributed computing infrastructure.

As with other infrastructure technologies, messaging is considered an enabling technology that might require considerable integration effort on all other components in the infrastructure. On the benefit's side, a deployed messaging infrastructure may be configured so that it can be simultaneously offered as a service to the customer in a dual-use model.

This capability may be potentially delivered through an ActiveMQ brokering network as it is currently used in parts of the EGI operational infrastructure (e.g. to deliver monitoring messages to a reporting service). In the long-term, EGI should consider a messaging infrastructure that is built on top of standards (such as AMQP [R 21]) may provide better sustainability options than the current solution. This discussion however will have to be discussed across all EGI and thus goes beyond the scope of this document.

5.1.4 Monitoring

It is important to monitor the current state of the infrastructure, as well as record the gathered measures for historic evaluation and prediction for future use as a means for Resource Providers for infrastructure capacity planning. The monitoring infrastructure will make use of the Messaging service to connect the monitoring emitters with the measures aggregation and reporting services.

This service is suitable for dual-use for infrastructure management and as a service to the customer (perhaps through the Notification service) in that VM state monitoring may be relayed to subscribed entities. Therefore, Monitoring must integrate with the federated AAI infrastructure outlined in section 5.1.1.

With SAM/Nagios, EGI already has a valuable asset in its portfolio for delivering this service. For improved sustainability standards based solutions should be preferred. However there are currently no standards available nor in the pipeline that would cover this capability,

5.1.5 Accounting

The purpose of Accounting is to monitor resource usage across the Infrastructure. While Monitoring is used to determine the current state of the Infrastructure, Accounting is used as a retrospective tool, even when the accounting interval is very short, e.g. 5 minutes. Typically, Accounting data is correlated with user information to be able to provide resource usage statements to customers. Therefore the Accounting solution must integrate with the Federated AAI infrastructure outlined in section 5.1.1. Therefore accounting data will be available per customer, and per operator – accounting data per individual end-user is neither necessary nor of interest to an Infrastructure Platform provider.

EGI already has a suitable accounting solution in its software portfolio. APEL has a proven history of collecting and aggregating accounting records in the European Grid community. APEL uses OGF Usage Records to convey compute-related accounting information. This will have to be extended for cloud-related accounting needs, and expanded to cover storage related accounting information as well. OGF currently develops a storage accounting extension to UR.

5.1.6 Information Discovery

The purpose of the Information Discovery service is to enable Research Communities to determine which of the federated Cloud Providers are suitable for deploying the respective Community Platform. This process is similar to choosing a mobile telephony service operator and subscription plan. In that sense the Information Discovery service may be seen as a comparison portal between the participating Cloud Infrastructure Providers.

Therefore the specific information that will be provided by this service ultimately depends on the EGI Research Communities' business needs. Since EGI is adopting the strategic goal of supporting many diverse user communities, this service must be able to serve many diverse inquiries coming out of the Research Communities. However, the type of information queried is not known, because an initial service that can be queried is not available – a chicken and egg problem.

Providing a service with an intentionally limited scope of provided information will start the virtuous cycle of continuous service improvement once Research Communities start using this service. Thus the starting set of information provided in this service will be:

- Registry of federated providers (through basic provider identification data). This allows querying for specific providers, and to identify matching providers in a result set after querying for other information.
- An indication whether a provider is currently accepting new customers. A Clout Provider may be part of the federation, but currently does not accept new customers. This allows the inquirer to filter out those providers.
- An indication of resources available to new customers, if applicable. This is a classification of available or "free" resources, not an accurate daily record of unused resources. This allows a Research Community to look for only those Cloud Providers that can single-handedly satisfy their resource requirements (if they do not want to scale across many providers), or with which Cloud Providers to engage in discussions if scaling across providers is necessary
- VM Endorsement policy indicator. Research Communities may not want to engage with Cloud Providers that enforce a strict and thorough endorsement policy (e.g. if the community cannot effort spending extensive effort on endorsement.
- VM housekeeping policy (i.e. enforced graceful shutdown of idle VMs). Infrastructure Providers expressed the concern that large numbers of idle Virtual Machines consume significant amounts of physical resources that may impede the provider's the overall performance in delivering services to the consumers. By implementing a policy of gracefully shutting down a VM, an Infrastructure Provider may safeguard the exposed virtual resources and make them available for other consumers.

The information is a combination of static and dynamic data. Automatic data should be taken in via the Message service, and provided by the Monitoring service. Static or semi-static information should be either accepted via an Admin interface accessible by federated resource managers, or also via the Messaging service. The service must provide a public access interface, both via web browsing and for automated enquiries through LDAP. No authentication and authorisation will be necessary for the public service.

The service must integrate with the federated AAI infrastructure to control which infrastructure management services are allowed to feed in update information.

BDII is a potential candidate for delivering this service.

5.2 EGI Collaboration Platform

The EGI Collaboration Platform facilitates EGI communities to collaborate with each other on top of the EGI Infrastructure Platform – those two platforms are complementary to each other, and will be provided by EGI.

This platform comprises of all necessary tools and services for collaboration across all EGI communities (see also [R 20]). It will enable sharing of Virtual Machines and data, provide services to manage data transfers within and across Research Communities, allow group membership management, and a number of social collaboration tools. These services should be delivered as centrally operated services (i.e. in a SaaS model), providing both a Web UI and an API.

This section describes types of services that are considered *potentially* useful to offer as a service to the EGI Research Communities. The specific solutions mentioned in each section are heavily based on assets already available in EGI, and do not claim to be exhaustive or comprehensive. For each service mentioned below – and any other service that will be discussed in the future – a comprehensive analysis of requirements, business models and available solutions will have to be conducted before a decision will be taken about including the service in the EGI portfolio or not. But that sort of discussion is outside the scope of this document.

5.2.1 Federated identity infrastructure

The federated identity infrastructure allows EGI communities to tap into already existing identity management systems of their choice without influencing, impeding or even compromising the identity management systems of sibling communities. Where the same system is chosen, collaboration and synergies may occur by using the EGI Collaboration Platform. However, it is important to note that this identity management system is implicitly independently managed from the identity management used in the Infrastructure Platform. They may overlap, be identical in choice of technology even, but there are no compulsory management ties between the two systems.

Several potential solutions are available without a conclusive decision. OpenID, Shibboleth and eduGAIN all provide similar functionality with small but distinctive differences.

5.2.2 Data movement

One key aspect of public e-Research is sharing data, from its production using all kinds of instruments to higher-level analysis and research paper publication in scientific journals.

Globus Online is a promising solution that provides an easy and lightweight data sharing management interface with distributed data transfer and access endpoints using several access and transport protocols, such as GridFTP.

An EGI-wide Globus Online service tied into the federated AAI for EGI communities allows EGI communities to share data without having to operate the necessary infrastructure beyond local data curation and control.

5.2.3 VM Image Sharing

The existing AppDB allows researchers to discover and share knowledge about existing scientific applications, facilitating re-use and reduction of unnecessary software development effort. However, application integration and packaging is still left to do, as an unnecessary barrier to scientists using ICT infrastructure for reusing existing software.

A symmetric service in the EGI Platform architecture allows Research Communities to discover and share VM Images comprising of scientific applications (or generic middleware services). This will significantly lower the barriers of researchers re-using existing applications - no more integrating and packaging of software before it can be used.

The StratusLab Marketplace provides a good starting point for a platform to share scientific applications packaged in Virtual Machines. Provided as a service, it integrates with independent appliance repositories that are managed locally. VMs may be stored in many locations, and identical copies are identified through unique computed identifiers. That way, more than one community may share the maintenance and provisioning of a set of appliances, and vouch for it by signing only one metadata entry in the marketplace.

EGI is currently running a closed test instance of an Appliance Repository and a VM Marketplace based on StratusLab technology.

5.2.4 Research group membership

Researchers taking part in large international and worldwide research projects are often members of more than one project. Multi-project affiliations require a lightweight and flexible infrastructure to join and leave research groups, particularly for short-term projects.

VOMS is a very popular tool to manage group membership information. Delivered as a service in the EGI Collaboration Platform, it may lower the necessary community specific cost of IT infrastructure thus lowering the barrier for new user communities to engage with EGI.

5.2.5 Service Desk

Efficiently operating a community platform requires a well-organised service desk for your users. Well-proven processes and tools exist that may be offered as software services to EGI communities, leveraging similarities in service desk operation and processes across user communities.

GGUS provides flexible deployment options for both global and local service desk functionality as a software service. It may be supplemented by live chat services based on XMPP/Jabber, public knowledge base services (either shared or delivered as an individualised service).

5.2.6 Meeting planning

Collaboration requires regular meetings. Whether organising phone conferences, focused Face-to-Face meetings or large conventions, conferences or community platforms, the requirements are almost always the same. Therefore, a globally accessible meeting planning service may help attracting new user communities.

Many different solutions exist; Indico is a very popular and well-known solution that used worldwide in the HEP community.

5.2.7 Training platform

Training is a ubiquitous need in a constantly evolving and renewing world. Generalising training as a means to pass on intellectual knowledge between individuals, a common training platform allows trainers and trainees to focus on the actual knowledge sharing.

Currently, EGI is providing a coordination function for training activities through the EGI Training Marketplace. However, this implies that the training provider already has at least a minimal training infrastructure in place before training services can be offered. Small Research Communities often do not have the resources or skills to establish a training infrastructure by themselves. By providing a training platform as a service to the EGI ecosystem, EGI may lower the barrier for those communities to offer training services.

Moodle is an e-Learning platform that integrates well with other services, such as Blogs, traditional websites, online forums etc. It is a generic e-Learning platform with a worldwide active development and user community. Moodle is frequently offered as an e-Learning platform for and by various communities worldwide, which may act as a blueprint for EGI to integrate Moodle as a common EGI ecosystem training platform.

6 COMMUNITY PLATFORMS

In support for current EGI Research Communities, and perhaps in the future for many more Research Communities, we expect that a multitude of different community platforms emerge that integrate with EGI's e-Infrastructure. From the EGI Infrastructure Platform perspective, Community Platforms form the business layer of a commonly known 3-tier architecture (see Figure **11**) by implementing the respective community's preferred way of using the virtualised resources. From a Community Platform user perspective, Community Platforms may also include the User access layer to the communitys e-Infrastructure, thus being a SaaS offering utilising EGI's Infrastructure Platform (i.e. a stacked platform as described in Figure 2.

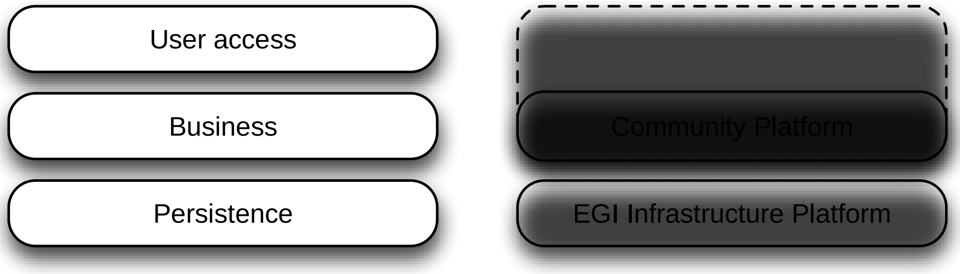

Figure 11: Classic 3-tiered architecture (left) compared to the EGI Platform model (right)

In that, individual Community Platforms form a higher-level infrastructure tailored to Research Communities. In other words they represent the community specific e-Research infrastructures federating geographically dispersed resources through the means of specific platforms.

6.1 Brokered HPC

Brokered HPC is an alternative approach to use distributed HPC resources. Users can see and use all the grid infrastructure as a whole. This approach has many advantages; heterogeneous grid resources can be used as a local HPC cluster. This mechanism is transparent for the end user, several grid resources are distributed in different places and sites but it acts as a local big HPC cluster.

The Polish scientific community (PL-Grid) is developing a brokered HPC solution based on QosCosGrid (QCG) [R 14] middleware. QCG is an integrated system offering advanced job and resource management capabilities to deliver to end-users a supercomputer-like performance and structure. User communities can execute a variety of applications, such as workflows, MPI or hybrid MPI-OpenMP applications over this layer. QCG can execute large-scale application written in Fortran, C, C++ or Java. These applications can be automatically distributed over a high-speed network of computing resources with guaranteed QoS. QCG also implements advanced features, such as resources reservation (similar to ticket and reservation system available on advanced batch systems like Torque or GE).

The Brokered HPC platform would be deployed in a hybrid setup. The HPC resources would continue being deployed in a traditional model since hardware virtualisation support is not available for critical components in the HPC model (e.g. for Myrinet networks). The higher-level Brokered HPC services would be deployed on top of the EGI Infrastructure Platform while utilising it for integrating HTC resources efficiently and seamlessly into the platform.

Another important feature of QCG is that supports open and standard based architectures (like OGF DRMAA, JSDL or BES); it uses secure communication channels (SSL/TLS, X.509) to authenticate each user and job.

6.2 Classic HPC

This platform represents the regular high-performance computing site running a high performance parallel environment. Through its architecture the Classic HPC Platform typically bypasses the EGI Infrastructure Platform.

The Classic HPC Platforms supports low latency networks (based on InfiniBand or Myrinet) to give the best performance possible to execute MPI jobs. MPI support was included into UMD repository to be used and installed by different user communities. Some of the MPI characteristics supported in EGI are:

- Parallel jobs can be executed by all the batch systems supported in EGI.
- Different MPI flavours are included like MPICH our OpenMPI.
- MPI information (cores availability, MPI flavours and versions..) is included and propagated by the grid information system.
- Different Nagios probes are executed to check MPI site sanity.

HPC sites can choose several infrastructure platforms like UNICORE. UNICORE (Uniform Interface to Computing Resources) is a ready-to-run grid service based on a client-server model. UNICORE supports various batch systems (Torque, SLURM, LSF, etc.) and provides clients for different operating systems (Windows, Linux, Mac OS X etc.). These features satisfy the needs of various scientific communities (e.g. graphical clients to define complex workflows, command line tool or web based access) and it can help to develop application integration in a HPC ecosystem.

6.3 Data-intensive HTC

This platform represents the current state of the art Grid middleware deployment that is predominant in EGI. Many of the services that are present in the current EGI production infrastructure are re-used in this description of the data-intensive Community Platform. Many if not all of the services of this platform require high-availability and load-balancing features, and are likely to consume large amounts of virtualised resources. All of these requirements can be satisfied by using the EGI Infrastructure Platform, given that some changes to the deployed services are implemented (e.g. the deployed service may be extended to access the IaaS management interfaces by itself to create a self-administering high-available service).

Existing security infrastructure services such as MyProxy, VOMS, and ARGUS can be easily deployed and managed on top of the EGI Infrastructure Platform with little or no modification for IaaS use. Traditional batch-queue orientated compute services have a large potential to utilise the EGI Infrastructure Platform as a replacement for Worker Node deployments (i.e. scientific applications that would be installed on traditional Worker Nodes will be encapsulated in Virtual Machines), with services such as WMS and CREAM changing from Batch Queue management systems to VM orchestration services. dCache, StoRM and DPM are typical storage solutions that satisfy different levels of storage demands. Already, some of these solutions essentially provide virtualised storage (in the sense of providing storage containers that expose access interfaces) that may be integrated as Storage Cloud management solution. Supplemental services such as metadata catalogues (e.g. LFC, AMGA), Service Registries (e.g. EMIR, UNICORE Registry) and general information discovery services (BDII, ARC InfoSys, etc.) are all suitable for deployment on top of virtualised resources since they do not have specialised requirements on compute resources.

6.4 Pilot-job HTC

Pilot jobs are widely used by e-Science Virtual Organizations for a long time already to execute their workloads. Most of these VOs are supporting different portals to provide all available grid resources to their users. For example, LHC experiments such as LHCb are developing and using the DIRAC portal to submit their pilot-jobs. These portals use pilot jobs instead of regular jobs in order to facilitate and extract the infrastructure complexity to their user communities. Users only need to connect to a specific portal using their personal certificate, the portal works as an intermediary between the underlying platforms (such as data-intensive HTC platforms) and users.

Pilot jobs are executed using the credentials of the portal that submits that initial, regular Grid job. During the time the pilot-job is executing, it continues fetching workload definitions from the workload server and executing them on the cluster using various security infrastructures, such as glExec. gLExec acts as a light-weight 'gatekeeper'. For each workload definition, glExec switches from the portal security context (with which the pilot job itself is executing) to a user security context that is primed with the identity of the user who submitted the workload definition to the workload server. In order to do so gLExec integrates with a number of local site security services, such as ARGUS, LCMAPS.

6.5 EGI Basic

EGI traditionally served, and still serves comparatively large Research Communities that make intensive use of distributed computing infrastructure. Being often referred to as "Heavy User Communities" these groups often have special requirements on their computing infrastructure. In contrast, many smaller research groups will have much less demanding requirements, and a carefully designed platform will satisfy many if not all their needs and potentially be simpler and easier to use.

Such a platform will be well integrated into the existing production infrastructure in terms of accounting, monitoring, information and general management (i.e. by integrating with APEL, SAM/Nagios, BDII etc.).

On top of these infrastructure needs the following may provide for most of the remote computing needs of smaller user communities. However, this platform pays much more attention to publically defined standards:

6.5.1 Compute

Computing jobs are formulated using JSDL, and submitted to OGSA-BES enabled services. Combined with a set of publically available extensions, such as the OGSA HPC-BP[5] or JSDL-SPMD[6], most common computing needs following the Job Submission paradigm should be satisfied.

EGI ha a number of components available satisfying these requirements. A combination of IGE supplied GRAM5 and GridSAM support MPI applications and the relevant standards (JSDL, BES, HPC-BP) enabling a rich client-side integration with this platform. UNICORE6 implements OGSA-BES and accepts jobs described in JSDL including the OGSA HPC-BP extensions. Both support a wide range of local resource management systems and batch queues such as Torque, LoadLeveler, LSF, and do not have demanding hardware requirements for deployment. Both technologies, Globus

[5] Open Grid Services Architecture – High Performance Computing Basic Profile
[6] JSDL Single Process, Multiple Data – a way to describe parallel applications.

and UNICORE, have shown over a decade substantial stability and are widely used in international Grid infrastructures (e.g. DEISA, TeraGrid/XSEDE, SkifGrid, etc.)

Some communities may require job-scheduling and management capabilities across management domains (e.g. when collaborating with more than one resource provider). GridWay, provided to EGI by IGE satisfies all those requirements.

6.5.2 Data/Storage access

Data needs to be accessible in a systematic manner, and the management and storage facilities need to grow with community need. Storage management and access will be exposed via popular standards, such as SRMv2, GridFTP, HTTP, and NFS4.

DPM is a lightweight storage management solution supporting all mentioned standardised access interfaces, has frugal resource requirements, and is easy to maintain.

7 OUTLOOK

The previous sections provided a technical starting point for EGI to start classifying and organising its existing assets (both IT Infrastructure and IT Services) in a way that enables collaboration and synergies based on independence and freedom, so that value added services may emerge and mature independently.

However, EGI has not arrived there yet. By beginning to deliver parts of their main assets, the large amounts of compute and storage capacities, Resource Infrastructure Providers federated with EGI are able to safely explore alternative means of software provisioning without degrading the agreed services to existing customers. By sharing and subscribing to the EGI Platform model, the first important steps are taken to provide a safe environment to enable *change*.

This change, though necessary, will not arrive in a big bang. This will include a change in ‘mind set’, ‘skill set’ and ‘tool set’ provided by EGI through the platforms approach in this document. Though steady and firm in its drive, change will have to be gradually implemented. This is in-line with the principle of nature leading to a steady growth through continues innovation. The EGI Federated Clouds Task Force is already working on realising all Capabilities of the EGI Infrastructure Platform as described in this document as a federated service delivery model, verifying its feasibility in a test-bed that is already used by early adopter Research Communities keen on seizing the opportunities that are offered.

As soon as key elements of the EGI Infrastructure Platform are considered stable enough for a federated deployment, committed Infrastructure Providers should be encouraged to deploy them for production use, and start taking on small focused user communities that require that particular capability while being in the position to wait for the others to stabilise. Naturally, the EGI Infrastructure Platform will have to mature somewhat ahead in time of the EGI Collaboration Platform and Community Platforms, since it is the enabling platform for the EGI Platform model as a whole. Therefore it will be of higher priority to provide a simple, yet complete set of services that comprise the EGI Infrastructure Platform. Once this initial set of services is deployed, it will allow EGI to start exploring the EGI Collaboration Platform. More importantly however EGI user communities then may begin integrating with the EGI Infrastructure Platform.

Existing Research Communities may gradually start integrating this platform into their current software and service portfolio, and begin to explore the feasibility of partially or totally migrating onto the EGI Infrastructure Platform. Feedback as to which features of existing services, and which type of services may be missing from the EGI Infrastructure Platform will be valuable material to discuss at the EGI Forum events held twice a year.

Deploying the EGI Infrastructure Platform provides new user communities with the opportunity to start using EGI’s e-Infrastructure services at the scale and flexibility that fits well with the community’s needs, thus significantly reducing the financial and resource barrier of integration. By simply re-using already available platforms (e.g. the EGI Basic Platform) in conjunction with the EGI training hub, the usage barrier is even more decreased keeping initial investments into the infrastructure low while exploring its capabilities.

7.1 Long-term future

It is very difficult to reliably describe what will, or even may happen in the long-term future. However, when comparing the EGI Platform model sketched in this document with other successful platforms such as Apple’s iPhone and iPad platform, it is perhaps not too far fetched to envision something similar to Apple’s App Store [R 15]: A platform almost of its own (though well integrated with existing platforms), providing a marketplace for appliances, all packaged to be deployed and operated at the expense of a couple of clicks: Compose your own operational platform by selecting what you need, and what is on offer. Adaptations of this principle already exist, for example the Ubuntu Software Centre [R 16]. Suitable precursors already exist in the EGI ecosystem – combine the EGI Applications Database [R 17] and StratusLab Marketplace [R 18], and the vision of an "EGI Platform Store" might become reality faster than one might expect.

8 CONCLUSION

EGI.eu, EGI-InSPIRE and the production EGI infrastructure are unique assets of the European research community well suited to satisfy the needs of a wide range of scientific disciplines such as all those thematic groups listed in the ESFRI roadmaps [R 1]. Reaching out to new user communities to expand the customer base is the central motivation for EGI to explore new ways how to deliver an efficient and easy to use trustworthy foundational infrastructure to its user communities and to evolve together with those to sustainable federated research infrastructures. By positioning itself as a ubiquitous federated e-Infrastructure within Europe, and a well connected and collaborating computing platform worldwide, EGI will prepare itself to support e-Research in the upcoming era of the European Research Area (ERA) [R 19] at a scale far in excess of its current activities.

This document describes an EGI Platform model and its initial architecture as a tool for EGI to reach this goal together with its user communities. EGI's current IT Infrastructure model of delivering a high-value end-to-end service to its current customers will continue to exist, with an EGI Infrastructure Platform that allows scaling out IT services across all new EGI Research Communities in a systematic way supplementing it.

By using a platform architecture, an infrastructure is neutral and impartial by definition in its support for its customers. Therefore the EGI Infrastructure Platform is designed to foster *choice* and *flexibility*, allowing for innovation and value-added services being built on top of it. Supplemented by the EGI Collaboration Platform, it will allow an interdisciplinary ecosystem to evolve on top of it that spans many research domains.

This approach will also allow Resource Infrastructure Providers federated in EGI themselves to reassess how they will deliver services to existing user communities: Either through continuing to deliver it in the traditional model, by transparently migrating it on top of the EGI Platform Model, or a mixture of both.

In turn, existing EGI Research Communities may assess which of the delivery models suit them best, and pick a choice. In fact, the heavy user communities of EGI may be seen as a blueprint for migration activities for many smaller user communities that may have fewer resources available to spearhead migration activities on their own.

The EGI Technical Roadmap [R 24] will provide more details on the further development of particularly the EGI Infrastructure Platform, and where possible, for the EGI Collaboration Platform, providing a comprehensive roadmap document across all technical activities within EGI. Currently being conducted based on best-effort contributions, the Federated Clouds Task Force and with the future development of the EGI Infrastructure Platform depends on the dedicated support of the NGIS.

9 REFERENCES

R 1	European Strategy Forum on Research Infrastructures, http://ec.europa.eu/research/**esfri/**
R 2	ITIL 2011 Glossary, http://www.itil-officialsite.com/InternationalActivities/ITILGlossaries_2.aspx
R 3	EGI Technology Glossary, https://wiki.egi.eu/wiki/Technology_Glossary_Contents
R 4	SearchServerVitualization, TechTarget, http://searchservervirtualization.techtarget.com/definition/platform
R 5	Oracle Technology Network, http://www.oracle.com/technetwork/index.html
R 6	Microsoft Developer Network, http://msdn.microsoft.com/en-us/ms348103
R 7	The Khan Academy, http://www.khanacademy.org/
R 8	MIT OpenCourseWare, http://ocw.mit.edu/index.htm
R 9	Edsger W. Dijkstra, Selected Writings on Computing: A Personal Perspective, Springer-Verlag, 1982. ISBN 0–387–90652–5.
R 10	Configuration Management Database (CMDB), http://www.cmdb.org/
R 11	IGI WNoDeS: Worker Nodes on Demand Service, http://web.infn.it/wnodes/index.php/wnodes
R 12	openModeller, http://openmodeller.sourceforge.net/
R 13	EGI Federated Clouds Task Force, https://wiki.egi.eu/wiki/Fedcloud-tf
R 14	QosCosGrid, www.qoscosgrid.org/
R 15	Apple AppStore, http://www.apple.com/mac/app-store/
R 16	Ubuntu Software Centre, http://www.ubuntu.com/ubuntu/features/ubuntu-software-centre
R 17	EGI Applications Database, http://appdb.egi.eu/
R 18	StratusLab Marketplace, http://marketplace.stratuslab.eu/
R 19	European Research Area, http://ec.europa.eu/research/era/index_en.htm
R 20	ESFRI project requirements for Pan-European e-infrastructure resources and facilities, https://documents.egi.eu/public/ShowDocument?docid=12

R 21	Advanced Message Queuing Protocol (AMQP), http://www.amqp.org
R 22	EGI Global Task review, MS115, https://documents.egi.eu/document/961
R 23	EGI Business Model, D2.18, to be published
R 24	D2.31: EGI Technical Roadmap, to be published

