



Post-EMI/IGE Software Provisioning

Collaborating with Technology Providers of varying levels of commitment

Michel Drescher (Michel.Drescher@egi.eu)

Steven Newhouse (Steven.Newhouse@egi.eu)

Gergely Sipos (Gergely.Sipos@egi.eu)

et al.

13 December 2012

Version 4

COPYRIGHT NOTICE

Copyright © 2012 EGI.eu. This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivs 3.0 Unported 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, USA.

The work must be attributed by attaching the following reference to the copied elements: “Copyright © EGI.eu (www.egi.eu). 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.

1 Introduction

With the EMI and IGE projects ending in Spring 2013, EGI.eu needs to plan and prepare for Software provisioning and collaboration with product teams and platform integrators that currently benefit from the coordination activities undertaken by EMI and IGE, identifying where EGI.eu can provide coordination activities and where these activities need to come from the Technology Provider or the User Community. The TCB members discussed this topic at the 14th TCB meeting [TCB-14], based on a proposal from EGI.eu [EMIPlan].

The transition to the proposed model includes services and coordination functions that are currently carried out by existing Technology Providers. A number of these services were discussed in [EMIPlan]. This document will discuss how these services will be continued and in which form: Some of these services will be covered by the TCB (section 2.1), some will be provided by Platform Integrators (section 2), and some will be part of ancillary service offerings tailored to Platform Integrators (section 5).

This document drills into the technical details of Software Provisioning, and how the various activities will contribute to, and be published in the UMD.

2 Technology Providers & Platforms

With the end of the EMI and IGE projects, EGI.eu is expecting a stronger partitioning of the Technology Provider (TP) landscape into a larger numbers of providers, each with stronger focus of interest around a particular capability or service. What's more, each of these providers will have their own sustainability plans and preferences for the level of commitment to collaboration and integration with EGI's provisioning processes. The EGI processes should meet these new expectations and should offer more options and allow more flexibility than they do today.

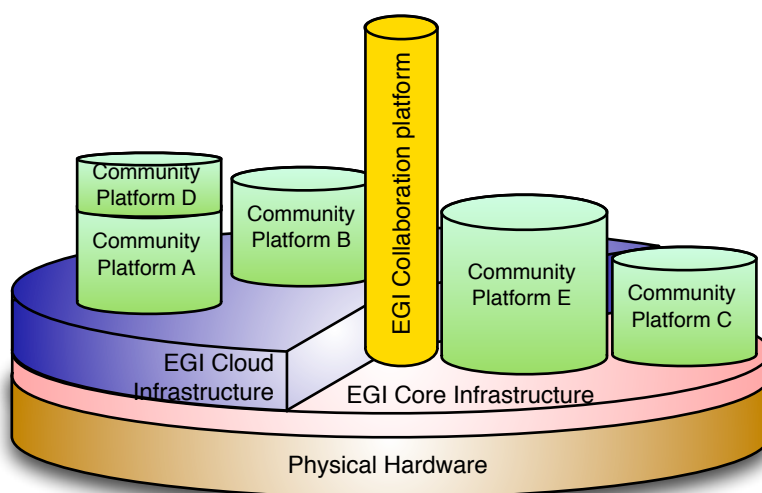


Figure 1: The EGI Platform model

In early 2012 EGI started developing a roadmap [MS510] that will gradually evolve the EGI production infrastructure into a platform oriented architecture, to be able to build an e-Infrastructure that is capable to support a broad customer base with a very diverse set of requirements. MS510 defines a number of platforms, and the associated stakeholders and actors in a



platform oriented architecture. Taken from MS510, **Error! Reference source not found.** gives an overview of the EGI Platform architecture.

The Resource Providers federated into EGI.eu's member NGIs own the physical hardware, the resources consumed by the research projects that are part of the EGI community. Those resources are federated together using the EGI Core Infrastructure Platform (services needed to operate a federation of locally deployed distributed computing platforms); technically this is achieved by Community Platforms (e.g. Platform E and C in **Error! Reference source not found.**) that are deployed directly on the physical hardware will have to integrate with number of services of the EGI Core Infrastructure Platform¹. The EGI Cloud Infrastructure Platform provides self-service provisioning and consumption of IaaS Cloud resources. Similar to Community Platforms, it is deployed directly on physical hardware, but as a generic platform it allows the deployment of any number of Community Platforms (e.g. Platforms A, B and D in **Error! Reference source not found.**) on top of it. The Core Infrastructure Platform also serves as a framework for the EGI Collaboration Platform in that it a number of components may integrate with the EGI Core Infrastructure Platform. The EGI Collaboration Platform is a set of generic and independent services, thus applicable to all supported research communities and are available to be consumed by humans, and/or by software that is part of any of the four platforms. Community Platforms are defined as providing infrastructure services tailored to the specific needs of the targeted EGI community.

EGI.eu will maintain ownership over the EGI Core, Cloud and Collaboration Platforms in terms of specification, integration and distribution where required. EGI.eu expects that Technology Providers will align their scope of activity with specific Community Platforms depending on their own expertise and the needs of the user communities that they may be associated with. The coordination of activities around the platforms deployed in EGI will take place at the TCB incorporating recommendations coming from domain specific management and coordination boards such as the OMB, UCB and others.

Therefore EGI needs to collaborate with Technology Providers on different levels. Identifying the different potential roles coming from the individual Technology Provides will help separating the requirements and responsibilities, as follows:

Platform Integrator (PI)

The role of a Platform Integrator is defined as defining the scope a specific platform serving its target community, and assembling a number of software components into regular releases of that platform². Applying this definition to the platforms defined above, EGI.eu will be its own Platform Integrator for the EGI Core Infrastructure, EGI Cloud Infrastructure and EGI Collaboration platforms. EGI.eu expects that external Technology Providers will assume the role of Platform Integrator for any number of Community Platforms. In particular, Platform . That said, an external Technology Provider may act as Platform Integrator for more than one Community Platform.

Platform Integrators are EGI.eu's main contact points of communication and coordination through the EGI Technology Coordination Board (TCB) in communicating technical requirements between the consuming users or deploying communities and the Technology Providers associated with a platform.

¹ This does not preclude Resource Provides deploying these services transparently in Virtual Machines. From an architectural perspective, this is an implementation detail.

² Naturally, this role includes ensuring that the included software and services provided by Product Teams work well with each other. If they don't, a Platform Integrator may replace an ill-behaving component with another, better fitting alternative.



Product Team (PT)

Product Teams produce software and/or provide well-defined services that are going to be included in all four platform types defined earlier in this document. Product Teams will work with the Platform Integrators of all platforms their products were chosen to be included in – it is a common use case that platforms re-use common or generic software components. For example, the CANL Product Team is likely to work closely with Platform Integrators that emerge from the EMI project and that decide to continue using CANL as an authentication library.

EGI.eu is not planning on working with Product Teams directly, except when fulfilling its role as Platform Integrator for the EGI Core Infrastructure, EGI Cloud Infrastructure and EGI Collaboration Infrastructure.

2.1 Changes to the TCB composition and relationships

The EGI Platform architecture will likely affect the composition of the TCB. A viable model is described in this subsection.

Clearly, Platform Integrators will be represented at the TCB, effectively substituting Technology Providers with a similar if not identical membership. Platform Integrators are representing the supply side of the technology relationships in EGI. Product Managers represent the demand side of this relationship. For each platform deployed in EGI one Product Manager will be the voice of the consuming customer at the TCB.

The activities carried out in the TCB will be divided into several committees governed by the general TCB assembly. These committees will be aligned with software release and provisioning activities as required for the EGI Platform architecture. Committees may convene independently from general TCB meetings focussing on their remitted activities, and reporting back to the TCB. The following committees incorporate in some form some services provided by current Technology Providers (see [EMIPlan]); more committees may be constituted as required:

- **Technical Coordination:** This committee will deal with the technical coordination, particularly the requirements management as defined in [TCBReqMgmt], and technical evolution discussions that currently take place in general TCB meetings. Integrated and Contributing Platform Integrators will be represented in this committee (as part of their commitment laid out in MoUs), meeting the Product Managers for the respective platforms. The Technical coordination committee will also produce and update Platform roadmaps at regular intervals.
- **Release Management/Coordination:** This committee – the UMD Release Team, URT (see below) – will coordinate the releases of Integrated and Contributing Platform Integrators (again as part of their commitment defined in MoUs). This coordination would be led by EGI.eu through weekly or bi-weekly meetings. Its purpose would be to collect upcoming release plans and provide the coordination necessary to ensure that new releases from Product Integrators do not break the existing contents of the UMD repository - a lightweight version of EMI's EMT meetings. EGI.eu effort would be small (0.5?) and would build on current resources used to do the existing UMD releases. The existing UMD release process would remain but EGI.eu would help resolve dependencies and technical conflict between PIs, and to a certain extent also PTs, in order to synchronize platform releases.

The general TCB assemblies are foreseen as a venue for regular reporting and governance body for the constituent focus committees. It will also deal with:

- **Sustainability and Exploitation:** EGI.eu’s focus will be on sustaining the integrated components and exploiting them for its own use within the Infrastructure Platforms. These services and software will be provided for others to reuse as they wish.

3 Technology Provider commitment levels

Recapturing the proposals made in [EMIPan], EGI.eu expects its Technology Providers to fall into three categories, integrated, contributing and community providers as summarised in Table 1 below.

TP type	MoU	SLA	TP QA	EGI.eu QA	URT	3 rd line support	EGI.eu benefits
Integrated	Y	Y	Y	Y	Y	Y	Y
Contributing	Y	Expected	Y	Equivalent	Y	Expected	Limited
Community	N	N	Optional	Optional	N	Optional	N

Table 1: High-level categorisation of EGI Technology Providers

The mechanics of the collaboration between Technology Providers of either type are described in Memoranda of Understandings (MoU) and accompanying Service Level Agreements (SLA) where applicable. These documents describe and define the scope of the collaboration, particularly the roles and responsibilities in the Quality Assurance (QA) of delivered software. The UMD Release Team (URT)’s role will be defined in these documents as the management body coordinating releases of software in the UMD as described below.

It is important to note that these terms – integrated, contributing, community – do *not* describe EGI.eu’s expectation of what type of platform a Technology Provider may provide. These terms rather describe the level of process integration and collaboration between EGI.eu and the specific Technology Provider.

However, EGI.eu will implement a policy that ties the criticality of a given platform to EGI.eu’s business success with the level of commitment of a Technology Provider to provide high quality software. This may translate into EGI.eu considering only Integrated Product Teams for services included in its Core Infrastructure and Cloud Infrastructure Platforms. Both Platform Integrators and Product Teams may collaborate with EGI.eu on either contributing or community level. It is up to the research community using a Community Platform to specify (through the appointed Product Managers) the quality and the corresponding QA process that it needs from the Technology Providers that comprise the respective platform.

Choosing a commitment level therefore has impact not only on how EGI.eu and the respective Technology Provider are collaborating, but also how the affected software will be received and perceived in the EGI community. Essentially, the relationship between Technology Provider and EGI.eu is governed by the higher the commitment of a Technology Provider, the higher EGI.eu’s commitment will be, too.



This includes not only the amount of benefits a Technology Provider may receive through other EGI.eu activities (e.g. Dissemination, priority access to slots in EGI community and technical fora), but also access to specific services EGI.eu provides in the area of Technology Provisioning.

4 Software Provisioning and (re)presentation in UMD

The Software Provisioning process needs to improve and adapt to the foreseen changes in EGI.eu's relationships with Technology Providers. This ultimately also affects the UMD, its contents, and how it is presented to its main customers, the Resource Providers federated into EGI.eu's member NGIs. The classification of Technology Providers into three categories (integrated, contributing, and community) will be reflected in the process of provisioning software into the UMD, where the integrated and community TPs form the two ends on a scale of Software Provisioning activities. Figure 2 provides an overview of the processes described in the following subsections.

4.1 Quality assurance and the provisioning process

The software provisioning process for Integrated Technology Providers will be identical to the process that is currently used. Integrated Technology Providers will conduct their own independent Quality Assurance and make software available for EGI to provision. The EGI.eu Software Provisioning teams will pull the software; it will verify it against its Quality Criteria and test in a Staged Rollout phase before it is made available in an integrated UMD "main" repository. The quality of the software will be monitored against the number of bug reports, post mortems of production infrastructure failures, vulnerability reports, and other KPI and KQI that are available for metrication.

The provisioning process for Contributing Technology Providers is expected to be similar to the process for Integrated Technology Providers except that the Technology Provider conducts the complete software quality assurance. Individual TPs may choose to use the existing EGI Software Provisioning process and tools, which EGI.eu is planning to offer as a service to Contributing Technology Providers. For those that do not use these services EGI.eu will audit the Quality Assurance documents, processes and artefacts on a regular basis to build its trust in the respective Technology Provider. Once this trust and agreement to operate as a contributing Technology Provider is in place, contributing Technology Providers will be given access to a tool that allows them to upload any number of software packages and corresponding release information into the EGI repository.

The provisioning process for Community Technology Providers is very simple, in that they will be allowed to upload any number of software packages to the repository at any point in time, without any expectation or constraints on quality assurance or even release timing formulated by EGI.eu.

Furthermore, the releases of integrated and contributing Technology Providers will be coordinated through regular meetings of the UMD Release Team, in which each Contributing Technology Provider will be represented along with the Integrated Technology Providers. This UMD Release Team will ensure that releases of Integrated Technology Providers are well coordinated, and that releases of the Contributing Technology Providers are synchronised accordingly to ensure a consistent set of software across Integrated and Contributing Technology Providers.

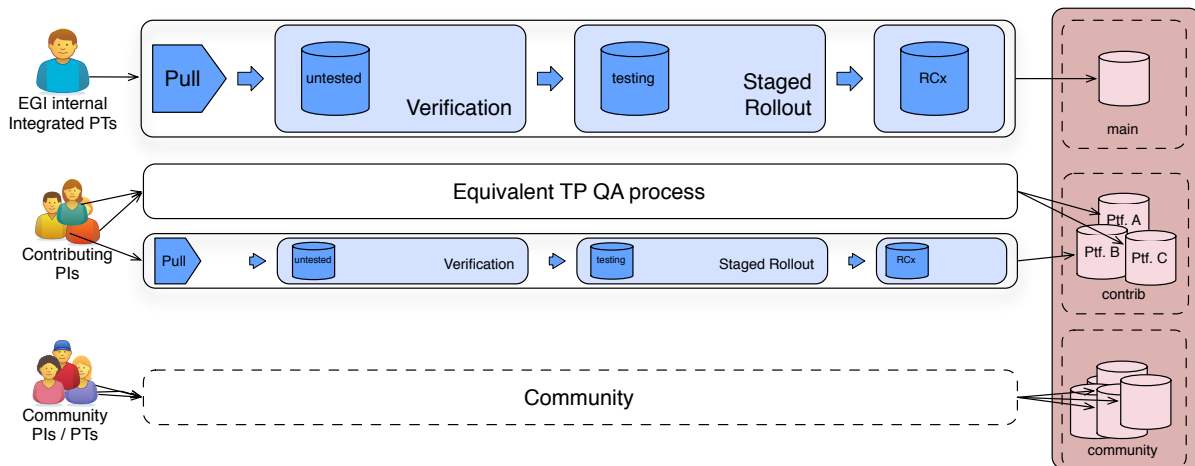


Figure 2: Overview of the UMD software provisioning

4.2 Changes to the UMD

The changes to the UMD are expected to be fairly significant in that it will no longer be a concise, integrated repository containing all software from all Technology Providers. Instead, the UMD may become a distribution of software that is deployed in EGI on various levels within the different EGI and community platforms.

These changes support a number of provisioning scenarios. For example, new Technology Providers may be included as Community Technology Providers (whether as Product Team or Platform Integrator) with a very lightweight intake procedure, determining the popularity of the provided software, before either side may consider a more formalised collaboration. Also, the functionality of the AppDB could be extended for any software entry to upload a release into the EGI Software Repository as a Community Technology Provider.

The following describes the technical layout of the repositories that will serve the Resource Providers with new and updated software. Figure 3 provides an overview of the layout of the UMD repository as a whole.

The UMD will be composed of three software *domains*, reflecting the three different levels of QA and release coordination. Each of these domains may contain arbitrary numbers of repositories³.

The *UMD main* domain will contain exactly one repository – the integrated UMD repository, as it exists today.

The *UMD contributed* domain will contain repositories providing software from Contributing Technology Providers. As a consequence, the UMD Contributed domain will contain repositories covering individual libraries, wrappers, up to complete Community Platforms⁴. Platform Integrators will assemble the contents of these platform repositories using components from the UMD contributing domain, UMD community domain, or other sources as they wish. The UMD Release

³ At this point, this document refers to repositories on the service level, not its actual technical manifestation. That is, the term repository refers to a means of providing software to the Resource Providers. The technical manifestation will be discussed elsewhere.

⁴ If a Platform Integrator wishes to maintain two Community Platforms, she will have to maintain two such repositories as well.

Team (a standing committee of the TCB, see section 2.1) will manage the coordination and synchronisation of the repository updates in this domain.

The *UMD community* domain will – similar to the UMD contributing domain – contain repositories providing software from Contributing and Community Technology Providers. This domain will also contain repositories of varying scope. However, unlike for the UMD Contributing domain, the updates of the various repositories in the UMD community domain will be *unmanaged*, i.e. Community Product Teams and Community Platform Integrators may update their individual repositories following their own release process and timing.

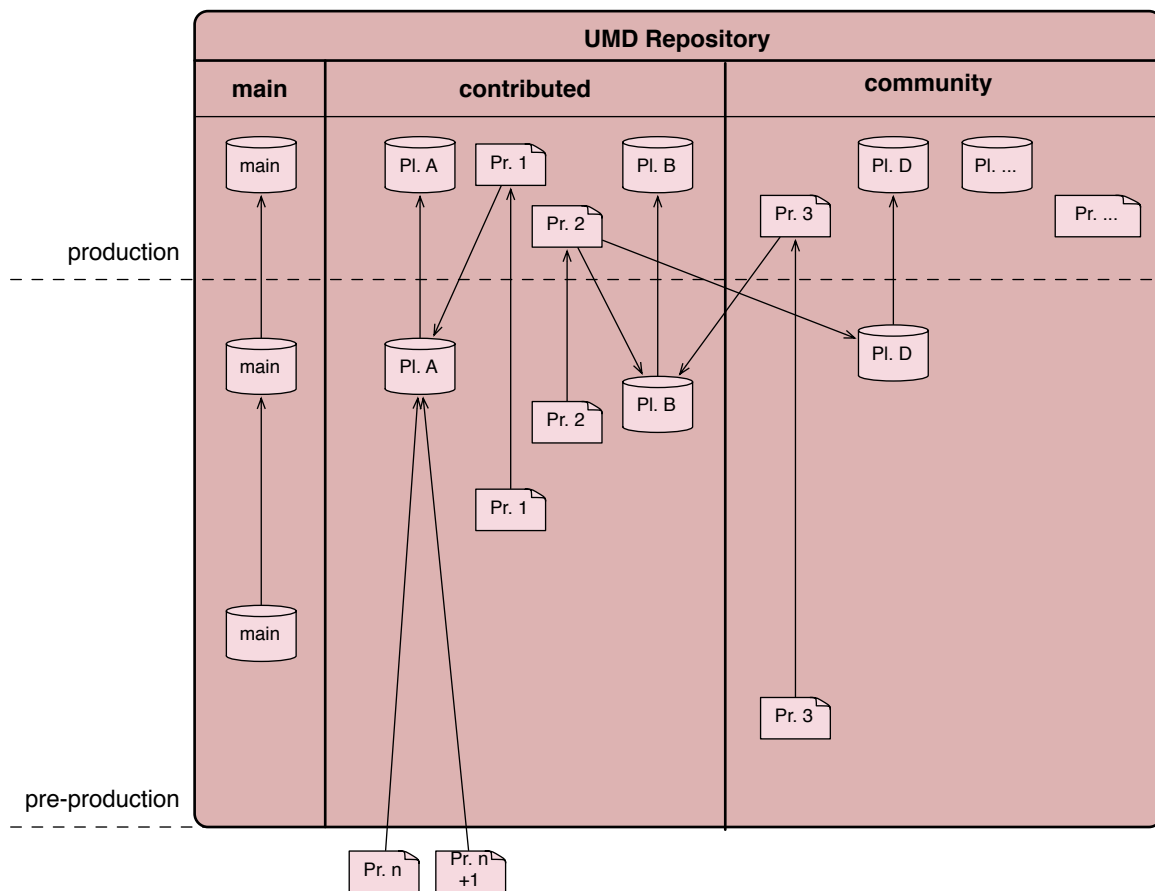


Figure 3: A conceptual view on the UMD repository

5 Service offerings

In order to implement the changes to the software provisioning processes and tools, a number of services will have to be made available to Technology Providers to be able to participate in this framework. The following subsections describe the various services available.

5.1 Services for Community Technology Providers

Community Technology Providers need very few services in order to participate in this software provisioning framework. Their entire work is uncoordinated, and releases are allowed to be made at any point in time in their respective repository.



The set of services for Community contributors is currently proposed as:

Basic package:

1. Access to one software repository in the UMD community domain.
2. Access to an uploading service; this service will support updating updated software for the repository, and a free-form text field (HTML text editor) allowing to provide release announcement information published at <http://repository.egi.eu> in an appropriate channel.
3. Common AAI of all services using EGI SSO (subject to availability).

Optional services:

1. One EGI Helpdesk support unit including access to basic reporting and statistics, without the integration into EGI's 1st level and 2nd level support processes.
2. One discussion forum at the EGI discussion forum service (<https://forum.egi.eu/>) for developers and users to connect and share. This forum will be in a "Community Software" group.
3. Access to the EGI blog facility.
4. Other, non-technical services, such as dissemination events, access to EGI community and technical for a may be available, but are out of scope for this document.

5.2 Services for Contributing Technology Providers

Contributing Technology Providers commit to and demonstrate considerable more collaboration and support to EGI.eu (and its federated members).

Technically, with regard to the UMD repositories, there are no differences to the requirements for Community Technology Providers provided that Contributing Technology Providers strategically align with the boundaries of Community Platforms (which EGI.eu desires).

The only exception to this is the alignment with major UMD versions: Contributing Technology Providers will receive access to as many individual software repositories as EGI is supporting major UMD versions. The current policy states that EGI will support major UMD versions for 2 years, with an overlap of one year between any 2 subsequent UMD major versions. This translates to two repositories for a Contributing Technology Provider, for any two UMD major versions being supported. If this policy will change in the future, the provisioning of repositories for Community Technology Providers will have to adapt as well.

EGI.eu expects from Contributing Technology Providers an end-to-end quality assurance programme that is equivalent to the programme undertaken together by Integrated Technology Providers and EGI.eu according to section 3. This QA programme is expected to be entirely covered by the respective technology provider's efforts and resources. However, there are EGI.eu services that cover essential tools and activities in a QA programme that Contributing Technology Providers may make use of, if they wish to do so.

The set of services for Contributing Technology Providers is currently proposed as:

Basic package:

All of the basic package for Community Technology Providers, except that software repositories will be aligned with major UMD releases (including the upload service), plus:

1. Software release coordination in a UMD Release Team (URT) together with Integrated Technology Providers.



2. Observation status in the Technology Coordination Board (TCB) (Platform Integrators only).

Optional services:

1. Access to a testing infrastructure where production versions of UMD main and Community Platforms of Contributing Platform Integrators are available to test against.
2. Voting member status in the TCB (an SLA must be in place and in force) (Platform Integrators only).
3. One or more GGUS support units (in case of more, one “head” SU will be included) including access to reporting and statistics, and integrated into EGI’s 1st level and 2nd level support processes.
4. Access to the EGI Software Provisioning process and tools.
5. One or more discussion fora (organised as sub-forums) at the EGI discussion board at <https://forum.egi.eu/>.
6. Access to the EGI blog facility.
7. Other, non-technical services, such as dissemination events, access to EGI community and technical for a may be available, but are out of scope for this document.

5.3 Services for Integrated Technology Providers

Integrated Technology Providers are not much different from Contributing Technology Providers. The main difference is that EGI will continue the TP’s software quality assurance process with its own software provisioning activities before the software will be accepted for inclusion in a new UMD release. Support through EGI’s Helpdesk system will be required for Integrated Technology Providers

Basic package:

All of the basic package for Contributing Technology Providers, except that EGI.eu will take the software through its Software Provisioning process and integrate it into a main UMD software repository, plus:

1. Voting member of the Technology Coordination Board.
2. Software will be validated against EGI Quality Criteria and exposed to pre-production in Staged Rollout through the EGI Software Provisioning process.
3. One or more GGUS support units (in case of more, one “head” SU will be included) including access to basic reporting and statistics, and integrated into EGI’s 1st level and 2nd level support processes.

Optional services:

1. One or more discussion fora (organised as sub-forums) at the EGI discussion board at <http://forum.egi.eu/>.
2. Access to the EGI blog facility.
3. Other, non-technical services, such as dissemination events, access to EGI community and technical events may be available, but are out of scope for this document.

6 References

[TCB-15]	15 th TCB meeting (F2F), 14 December 2012, http://go.egi.eu/TCB-15
[TCB-14]	14 th TCB meeting (F2F), 6 November 2012, Amsterdam, NL, http://go.egi.eu/TCB-14

[EMIPlan]	Plan around EMI - https://indico.egi.eu/indico/getFile.py/access?sessionId=5&resId=0&materialId=1&confId=1170
[MS510]	MS510: EGI Platform Roadmap, https://documents.egi.eu/document/970
[TCBReqMgmt]	TCB Requirements Management process, https://documents.egi.eu/document/440

Appendix A Provisional platform compositions

This appendix defines a provisional list of components included in the platforms that EGI is retaining ownership of. Technology Providers may extend this overview with any number of Community Platforms on top of this in subsequent editions.

A.1 EGI Core Infrastructure Platform

The EGI Core Infrastructure Platform (Core Platform for short) comprises of operational services that support the operational management of a federated infrastructure, and a number of technical software services that aid in integrating any other platform with the operational management infrastructure. The following list includes existing services and, where applicable potential additions (indicated by a question mark in brackets: (?))

Operational Services:

1. Operations Portal
2. GGUS (EGI Helpdesk)
3. Metrics Portal
4. Gstat

Technical Services:

1. Messaging infrastructure
2. SAM (Monitoring)
3. GOCDB (Information)
4. APEL (Accounting)
5. AAI
 1. EGI trust anchors (EUGridPMA)
 2. VOMS (Attribute Authority for DTEAM, OPS)
 3. Gridmap files, etc.
 4. ARGUS?

A.2 EGI Cloud Infrastructure Platform

The EGI Cloud Infrastructure Platform (Cloud Platform for short) comprises of locally managed IaaS Cloud services, and a number of collaborative tools that may in the future become part of the Collaborative Platform. The main IaaS Cloud services are Computing and Storage services; Resource Providers deploy and maintain appropriate Cloud service components as long as these are integrated with the Core Infrastructure and expose the Compute and Storage services through standardised interfaces.



IaaS Cloud service stacks:

1. OpenStack (Open Source)
2. OpenNebula (OpenSource)
3. StratusLab (Open Source extension of OpenNebula)
4. ~okeanos (GRNET)
5. WNoDeS (IGI)

Integrations and others:

1. rOCCI server – provides integration with the Core Infrastructure AAI (GWDG et al)
2. rOCCI client – OCCI command line client and API (GWDG et al)
3. APEL integration for OpenStack
4. APEL integration for OpenNebula
5. Nagios plugin for OCCI Cloud Compute services
6. Nagios plugin for CDMI Cloud Storage services
7. VOMS integration for EGI VM Appliance repository

Collaborative tools and higher level services:

1. VM Marketplace (StratusLab component)
2. Vmcaster, in consideration (DESY, soon private engagement by Owen Synge)
3. Brokers, several under consideration
4. Portal software, WS-PGRADE and INFN portal under investigation.

A.3 EGI Collaboration Platform

Provisional list:

1. Federated Identity Management infrastructure
2. Data movement services (e.g. Globus Online)
3. VM Image Sharing (likely to take over VM Marketplace and VM Caster)
4. Research group membership (e.g. VOMS) for VOs
5. EGI Service Desk (currently GGUS)
6. Meeting planning (Currently Indico)
7. Training Marketplace
8. Application Database
9. CRM system

Appendix B Technical architecture of UMD repositories

By design, section 4.2 describes repositories on a conceptual level, treating them as services with unspecified technical realisation. This is to not take the reader's attention away from the higher-level concept that EGI.eu is planning to implement and expect Technology Providers to find their place in.



This section discusses in more detail possible technical manifestations of these repositories and the implications on how the UMD service is offered to Resource Providers, who will have to enable repositories on their physical infrastructure in order to support a number of research communities.

From a technical point of view, two types of repository manifestations are foreseen to coexist in the UMD infrastructure:

- a) VM image repositories, and
- b) Package management repositories, typically integrated with APT based package management (DEB packages for Debian and derivative Oss) or YUM based package management (RPM packages for Red Hat Linux and derivatives, SuSE Linux, etc.)

Either technical manifestation is treated as a “first-class citizen” in UMD. However, there are pros and cons to both approaches, making either more or less suitable for Platform Integrators and Product Teams.

VM image repositories are usually simple storage locations providing a download link for each contained VM image with little or no database/metadata overhead attached. VM images are Operating System independent in that they incorporate an OS within them as a bootable component. However, VM images require the support of specific hypervisors, leading to the same appliance encoded into one or more VM images for the different supported hypervisors.

Package management oriented repositories contain packages (containers of data with inlined metadata) that can depend on other packages. Each package contains a certain defined functionality. The dependencies encoded in the packages form a directed acyclic graph, that all together deliver a specified set of functionality ranging from library level functionality up to a complete Platform. Packages in that sort of repositories are usually provided as binary packages, i.e. containing binary software elements compiled against a specific OS and hardware architecture. Depending on the compiler settings and the actual software design, binary packages may apply to more than one OS and hardware architectures at the same time. Nonetheless, package oriented repositories are most often organised following certain OS/Architecture combinations: The current UMD repository is no exception in that it provides (most of) the same software in three different technical repositories:

- Scientific Linux 5 on 64 bit hardware (SL5/64 bit)
- Scientific Linux 6 on 64 bit hardware (SL6/64 bit)
- Debian Squeeze (Debian 6) on 64 bit hardware (Debian6/64 bit)

Technically, the dependency trees of packages may span more than one package repository, so that the system administrator in most cases will have to enable more than one repository to be able to install the required software to provide a Community Platform for a supported research community. Again, UMD is no different in that system administrators will have to enable two, mostly three, repositories to be able to install, for example, a CREAM CE:

- a) The Operating System repository (e.g. for Scientific Linux 6)
- b) The EPEL repository (an add on repository) for SL6
- c) UMD

This is also true for system administrators using the EMI repositories directly.

This makes package oriented repositories more suitable for Product Teams (exceptions will prove this rule), while Platform Integrators may use package oriented repositories or VM image repositories.



While the technical implications for VM image repositories are very clear in that each VM image is physically self-contained, this is not true for package oriented repositories (see above). This allows Platform Integrators effectively to follow essentially two platform delivery models, each with different implications for the Resource Providers that must enable these repositories to deploy a Community Platform:

- a) Inclusion by reference, and
- b) Inclusion by copy.

The inclusion by reference model avoids duplication of binary packages in a repository since dependencies of one package to another may only be satisfied by enabling another repository. This model aligns well with scenarios where individual repositories are managed in different authoritative domains, but encourages platform fragmentation and “repository explosion” in distributed deployment scenarios for software federations.

The inclusion by copy model draws a tight border around the respective platform in that it includes physical copies of the required packages in one location – the repository becomes self-contained and is very similar to a VM image deployment. Ideally, the system administrator needs to enable only one repository in order to support a specific Community Platform.

From an overall operational perspective, the inclusion by copy platform delivery model is preferred over the inclusion by reference model in that it transports a clearer message in the EGI community (at the expense of somewhat higher assembly effort on the Platform Integrator side). EGI.eu will not prescribe either delivery method (or variations), but require instead that Platform Integrators will clearly document the platform delivery method in their overall community platform system administrator documentation.