



EGI-InSPIRE

EGI TECHNICAL ROADMAP

EU DELIVERABLE: D2.33

Document identifier:	EGI-InSPIRE-D2.33-FINAL
Date:	04/06/2013
Activity:	NA2
Lead Partner:	EGI.eu
Document Status:	FINAL
Dissemination Level:	PUBLIC
Document Link:	https://documents.egi.eu/document/1706

Abstract

This deliverable constitutes the second edition of the EGI Technical Roadmap. It provides an executive overview about the project's planned activities for project year 4. Where applicable, interactions and collaborations with external projects are briefly mentioned.



I. COPYRIGHT NOTICE

Copyright © Members of the EGI-InSPIRE Collaboration, 2010-2014. See www.egi.eu for details of the EGI-InSPIRE project and the collaboration. EGI-InSPIRE (“European Grid Initiative: Integrated Sustainable Pan-European Infrastructure for Researchers in Europe”) is a project co-funded by the European Commission as an Integrated Infrastructure Initiative within the 7th Framework Programme. EGI-InSPIRE began in May 2010 and will run for 4 years. This work is licensed under the Creative Commons Attribution-Noncommercial 3.0 License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc/3.0/> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, and USA. The work must be attributed by attaching the following reference to the copied elements: “Copyright © Members of the EGI-InSPIRE Collaboration, 2010-2014. See www.egi.eu for details of the EGI-InSPIRE project and the collaboration”. Using this document in a way and/or for purposes not foreseen in the license requires the prior written permission of the copyright holders. The information contained in this document represents the views of the copyright holders as of the date such views are published.

II. DELIVERY SLIP

	Name	Partner/Activity	Date
From	Michel Drescher	EGI.eu / NA2	02/05/2013
Reviewed by	AMB & PMB		24/05/2013
Approved by	AMB & PMB		05/06/2013

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	2 May 2013	Initial internal version	Michel Drescher, EGI.eu
2	24 May 2013	Second draft for internal review	Michel Drescher, EGI.eu
3	30 May 2013	Final draft incorporating review comments	Michel Drescher, EGI.eu Steven Newhouse, EGI.eu, Diego Scardaci, INFN

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/>.



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 deliverable constitutes the second edition of the EGI Technical Roadmap. It provides an executive overview about the project's planned activities for project year 4. Where applicable, interactions and collaborations with external projects are briefly mentioned.

The first edition of the EGI Technical Roadmap begun structuring its content around the three pillars of the EGI strategy, namely: (1) Operational Infrastructure, (2) Community & Coordination, and (3) Virtual Research Environments. This second edition continues with this structuring, and taking it a step further by putting a multi-purpose operational infrastructure that is fit for serving the requirements of H2020 at the heart of the EGI technical roadmap. This is not to diminish the importance of the other strategic activities: without these complementary pillars, EGI would not be able to deliver an operational infrastructure at its best potential. As outlined in the EGI Platform Roadmap the evolution of the EGI operational infrastructure into a distinct set of building blocks – three EGI-owned infrastructure platforms on top of which research communities deploy their community platforms as required – allows a clearer definition of the supporting activities and their purpose around the EGI ecosystem.

The **Core Infrastructure Platform** provides services that are required to provide management and uniform delivery of services in a federated environment. The general direction of the development can be summarised as automation and sustainability. Many activities aim to provide automated implementations of processes and actions that were once implemented and defined as manual, human-driven policies. Sustainability activities address the issue of continued funding for maintenance and development of the current solutions; typically this involves customer diversification.

The **Cloud Infrastructure Platform** is re-using EGI's decade-long experience of federating distributed infrastructures by strategically integrating with the Core Infrastructure Platform. By providing and mandating a thin federation layer, participating Resource Providers maintain a high degree of freedom in their own sustainability choices and deployment of Cloud Management Frameworks while taking up the increased responsibility of maintaining interoperability with the mandated interface support that is part of the federation layer.

The **Collaboration Platform** emerges to provide services that can be used across all platforms that are deployed in the EGI production infrastructure. As such it begins to position itself and the included services as the “service integration glue” collating EGI platforms, Community Platforms and other services into consistent technical virtual research environments.

Community Platforms are part of the Virtual Research Environments used by EGI's researchers, but are deployed remotely as part of EGI's Operational Infrastructure. Community Platforms emerge to exist in two “flavours” with two similar, yet distinct business models attached to them: Typically, Community Platforms directly integrating with the Core Infrastructure Platform¹ are deployed and operated by EGI (as distinct service to the target research community), whereas Community Platforms that integrate with the EGI Cloud Infrastructure Platform follow one of the characteristic Cloud computing principle of “self-serving”, i.e. deployment and operation of the Virtual Machines and services packaged within are provided by the Research Community itself. Originating from the software delivered by the EMI and IGE project, the number of Community Platforms is beginning to diversify and will continue to do so in PY4.

The inclusion of NGI International Liaison contacts, EGI Champions and virtual teams have proven to provide an effective method to channel and direct communications, and to find a platform of technical collaboration where common need is apparent. The virtual teams model has been expanded for PY4, by introducing dedicated funds through mini projects. These provide a model to focus funding on the

¹ This situation resembles very much the legacy deployment and operations model in EGI of the past 10 years.



strategic improvement of existing solutions with a clear goal of insertion the results of the successful projects into immediate day-to-day operations.



TABLE OF CONTENTS

1 INTRODUCTION	8
2 OPERATIONAL INFRASTRUCTURE	10
2.1 Core Infrastructure Platform	10
2.1.1 Technical Services	11
2.1.1.1 Federated AAI	11
2.1.1.2 Information Discovery	11
2.1.1.3 Monitoring	12
2.1.1.4 Accounting	12
2.1.1.5 Messaging	13
2.1.1.6 Central Service Registry	13
2.1.2 Support Services & Tools	14
2.1.2.1 Operational security	14
2.1.2.2 Operations Portal	15
2.1.2.3 Metrics portal	16
2.1.3 Software Provisioning	16
2.1.3.1 Quality Assurance	17
2.1.3.2 Quality Control	17
2.1.3.3 Staged Rollout	18
2.2 Cloud Infrastructure Platform	18
2.2.1 Federated Cloud management services	18
2.2.1.1 VM Management	18
2.2.1.2 Data Management	19
2.2.1.3 Image management and distribution	20
2.2.1.4 Information Discovery	21
2.2.2 Core Infrastructure integration	21
2.2.2.1 Virtual Organisation management & AAI	21
2.2.2.2 Monitoring	22
2.2.2.3 Accounting	22
2.2.2.4 Central Service Catalogue	22
2.3 Collaboration Platform	23
2.3.1 Technical collaboration services	23
2.3.1.1 Service Desk	23
2.3.1.2 VM Marketplace & Appliance Repository	24
2.3.1.3 Software Repository	24
2.3.1.4 Application Database	24
2.3.2 Social collaboration services	25
2.3.2.1 Training Marketplace	25
2.3.2.2 Client Relationship Management	25
3 COMMUNITY & COORDINATION	27
3.1 Human Networks	27
3.1.1 NGI International Liaisons	28
3.1.2 EGI Champions	28
3.1.3 Operations network	29
3.1.4 Geek Squad	29
3.1.5 Virtual Teams & mini Projects	29
3.2 Communications and marketing	30
3.2.1 Major EGI Events & Workshops	30



3.2.2	Materials.....	30
3.2.3	Main website and web design.....	31
3.2.4	Publications.....	31
3.2.5	Dispatches.....	31
3.2.6	External Publications.....	31
3.2.7	Media and public relations.....	31
3.2.8	Social Media.....	32
3.3	Strategy and Policy.....	32
3.3.1	Strategic Analysis.....	32
3.3.2	Service Management.....	33
3.3.3	Business Development.....	33
3.3.4	Policy Development and Support.....	34
3.3.5	EGI Compendium.....	34
4	VIRTUAL RESEARCH ENVIRONMENTS.....	35
4.1	Community Platforms.....	35
4.1.1	Unified Middleware Distribution (UMD).....	36
4.1.2	Emerging Community Platforms.....	36
4.2	Support Services for VREs.....	37
5	CONCLUSIONS.....	39
6	REFERENCES.....	40

1 INTRODUCTION

Supported through the EGI-InSPIRE project, EGI focuses on delivering a production-quality e-Infrastructure to the supported research communities. EGI has identified a number of activities that together are necessary to achieve this goal. The EGI Strategy groups these activities into “the three pillars of the EGI vision”, summarised as follows:

Pillar 1 – Operational Infrastructure: The Operational Infrastructures provides the technical ICT foundation of the EGI e-Infrastructure by providing a distributed, federated service platform for access by end users. Depending on the needs of the targeted research community, EGI offers federation and operational services, Cloud Infrastructure services, or collaboration services.

Pillar 2 – Community and Coordination: Integral to delivering a pan-European e-Infrastructure are services around social aspects of a large and complex e-Infrastructure. What is often described as “connecting people” includes community building, development of human capital, coordination, communication, and last but not least strategy and policy related activities across the entire EGI community.

Pillar 3 – Virtual Research Environments: Virtual Research Environments (VRE) are defined as the complete and inclusive work environment that is owned, deployed, managed and used by one or more closely related research communities. This definition includes ICT resources that are entirely remote and external to EGI as well as EGI resources that are, or will be, integrated into potential VREs. Support for Virtual Research Environments includes infrastructure services such as deployment and hosting of Community Platforms on top of EGI resources, but also consultancy and technical services for existing and new community services.

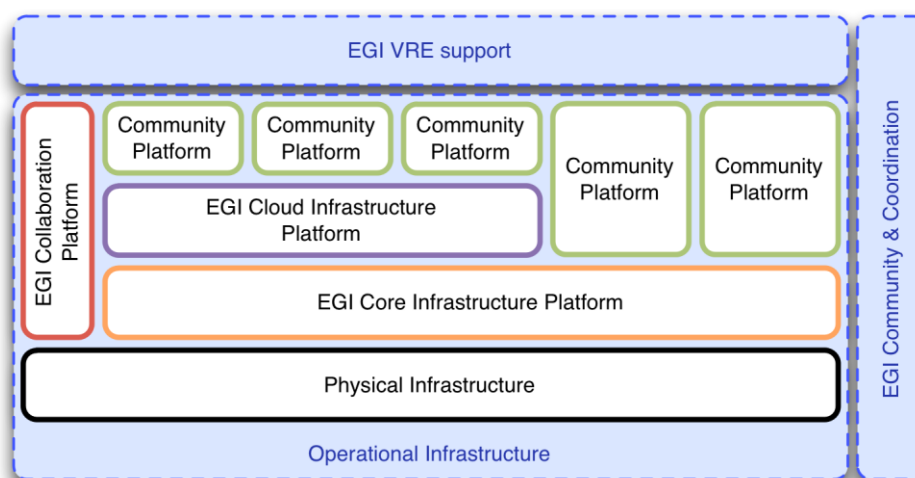


Figure 1: The three pillars of the EGI strategy

Receiving annual updates, the EGI Technical Roadmap is based on its first edition and its predecessor, the DCI Collaborative Roadmap. However, structure and content of this second edition reflects the changes and advancements of the past year. In a nutshell EGI is, technically speaking, transitioning into a service oriented business architecture, covering both the technical architecture and the business model. As a consequence, the deliverables and milestone documents on which this document is based in fact spread over a several sections and subsections in this document. For example, D5.9 contains material that is used in sections 2.1.3 and 2.3.1.3, and D7.3 together with MS710 contribute material to the sections 2.1.1, 2.1.2 and 2.3.1.



In and by itself this roadmap does not contain detailed technical development plans. Rather, it provides an executive overview, highlighting key developments and putting them into a consistent context across the whole EGI ecosystem. Information was sourced from a variety of technical documentation, and grouped in summarised form according to the structure indicated below.

Consequently, the EGI Technical Roadmap is structured around the three pillars supporting EGI's strategy towards H2020 as follows. Section 2 describes the planned improvements spanning the operational infrastructure. Section 3 describes the technical enhancements and plans necessary to deliver community and coordination services. Section 4 completes the roadmap with summarising EGI's technical support for Virtual Research Environments and Research Communities. The document ends with conclusions provided in section 5.

2 OPERATIONAL INFRASTRUCTURE

The operational infrastructure is organised into three distinct platforms that are for the most part owned and operated by EGI or, in case of some services that are part of the Collaboration Platform, by selected and trusted external partners.

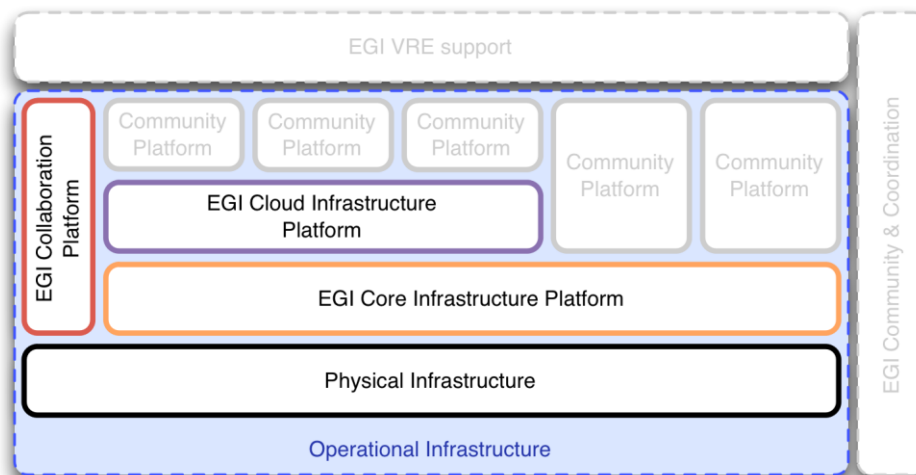


Figure 2: The second pillar of the EGI Strategy is organised in distinct platforms

Although Community Platforms are technically included in the EGI Operational Infrastructure as depicted in Figure 1 and in Figure 2 above, they are not part of EGI’s service portfolio. The reason lies in the ownership and offered services around these platforms. While EGI clearly owns and operates its Core Infrastructure Platform, Cloud Infrastructure Platform and (with some exceptions) the Collaboration Platform, this is not the case with Community Platforms: These are owned (defined, assembled, deployed and maintained) by the respective consuming Research Communities. EGI may collaborate with the Research Community to deploy and operate these platforms on top of its Core Infrastructure Platform (as is currently the case with WLCG and the UMD). In the future a ‘self-service’ model built around the Cloud Infrastructure Platform will allow all communities to deploy the Community Platforms they need as part of their Virtual Research Environment.

This section covers the technical activities planned for the Core Infrastructure Platform, the Cloud Infrastructure Platform and the Collaboration Platform. Section 4.1 provides an overview of Community Platforms at the time of writing of this document.

2.1 Core Infrastructure Platform

The Core Infrastructure Platform is scoped to provide “*Operational services necessary for the management of federated DCDis (Distributed Computing and Data Infrastructures)*”. Integrating and re-using a number of services from the Collaboration Platform such as the Service Desk, the EGI Software Repository and the Requirements Tracker, this represents the current operational infrastructure.

The Core Infrastructure Platform consists of a number of technical services; some of these are mandatory integration targets for any other platform deployed in the EGI production infrastructure (including the EGI Cloud Infrastructure Platform).



The Core Infrastructure Platform is completed by a number of support services and tools (mostly dashboards and portals) that are for exclusive use of the operations personnel in EGI.

Although technically part of the operational support services, the Software Provisioning activity is pivotal to ensuring a controlled and coordinated stream of updates to the various Community Platforms that are deployed in EGI. Therefore an entire subsection is devoted to this process.

2.1.1 Technical Services

2.1.1.1 Federated AAI

EGI does not directly provide its own Authentication infrastructure. Instead, EGI is collaborating with the EUGridPMA initiative, which in turn collaborates with the International Grid Trust Federation for policies and requirements around using X.509 certificates as authentication tokens for end-users and Grid services. Practically, EGI adopts the EUGridPMA's policies on trusting national academic certification authorities (usually provided by NRENs), and re-endorsing these in an extended form and scope, as EGI's trust anchor policy.

EGI-wide support for SHA-2 security algorithm: Driven by the SHA-1 security algorithm being considered a weak algorithm, EGI will transition globally to using certificates that employ the successor algorithm SHA-2. The transition will be mainly handled by operational security support (see section 2.1.2.1

2.1.1.2 Information Discovery

In PY3 most time was spent on developing and releasing the EGI GLUE2 Profile², and continuous upgrade of Grid components to versions that publish GLUE2 information according to this profile. However, this work has not been completed, and will thus continue in PY4:

Upgrade to EMI-2 and EMI-3 Community Platforms: Resource Centres will continue upgrading from old gLite 3.1, gLite 3.2 and EMI-1 deployments to EMI-2 and, preferably, to EMI-3 Community Platform deployments for ubiquitous EGI GLUE2 Profile support. This will include both GLUE2 publishing services as well as consuming services such as WMS, GFAL and ginfo.

Promote deployment of EMI-3 top BDII service: The EMI-3 top BDII service allows end-to-end publication of endpoint URLs expressed as GLUE2 LDAP URLs across GOCDB, top-BDII and site-BDII. This will also include exploring a transition path to including GSTAT 3 once it supports GLUE2.

Finalise EGI GLUE2 Profile using validator service: With such a large number of services coming from the EMI-2 and EMI-3 Community Platforms supporting the EGI GLUE2 Profile deployed almost at once, we expect a number of inconsistencies and bugs in the implementations. A central validator service will help removing these into a consistent EGI GLUE2 Profile support across the production infrastructure.

² <https://documents.egi.eu/document/1324>



Evaluate EMIR as global distributed dynamic service registry: The EMI Registry (EMIR) is a new component of the EMI-3 Community Platform that may serve use cases that do not entirely fit the Central Service Registry (see section 2.1.1.6) nor the current Information Discovery subsystem.

Include GPGPU resources: General Purpose Graphics Processing Units (GPGPUs) are increasingly used as computing resources for a number of use cases. To be able to integrate GPGPU resources into the EGI production infrastructure, the GLUE2 standard needs to be extended (and naturally incorporated into the EGI GLUE2 Profile).

2.1.1.3 Monitoring

The Service Availability Monitoring (SAM) component has accomplished most of the foreseen roadmap over the past EGI-InSPIRE project years. For the upcoming year, the following is planned:

Nagios probes for EMI-2 and EMI-3 Community Platforms with SAM Update 22: SAM Update 22 is expected to be published in PQ13, which will contain full support for EMI. Further probes for deployed EMI services will be integrated into future SAM releases.

Automated OLA violation reporting: In collaboration with the Operations Portal (see section 2.1.2.2) underperforming services will be automatically reported to the OMB for monthly review.

Support for on-demand monitoring: A number of day-to-day operational duties require much more flexible and dynamic deployment of monitoring probes than the current deployment model (probes bundled with SAM updates) is able to support. Such on-demand probes will be developed according to requirements defined by the OMB.

2.1.1.4 Accounting

EGI's Accounting service is a distributed infrastructure on its own account. It consists of a publicly available *Accounting Portal* and a central *Accounting Repository*. Traditionally, it was complemented by a client (the APEL client, developed and maintained within the EMI project) bridging the proprietary accounting record transmission protocol. The APEL client had to be deployed at the federated Resource Centres in order to provide accounting records.

The following new features are planned for introduction over the course of PY4:

Regional APEL repositories: Regional aggregation and/or distribution of accounting data is a desirable feature in a federated distributed infrastructure. Regional instances of the central APEL service allow independent processing of locally scoped accounting data as well as scaling out the Accounting infrastructure.

Messaging integration: The components of the Accounting infrastructure will be further modularised, reducing the dependencies between the components, and improving the flexibility in upgrading and deployment of modules in the infrastructure. The major target of the modularisation is the integration with the EGI Messaging infrastructure through the introduction of the Secure STOMP Messaging (SSM) protocol between the various Accounting subsystems. This includes a generic, lightweight SSMv2 client as a convenience tool for Community Platform integrators. The EGI Cloud Infrastructure Platform makes use of this lightweight client for integration purposes.



Extending accounting data coverage: Traditionally, Accounting in EGI focused mainly on tracking the usage of compute resources, such as CPUs, RAM and wallclock time. With EGI taking on the challenge of H2020 and positioning itself for the ERA, accounting of more resource types becomes a necessity. Extended accounting support will be implemented end-to-end, from the emitter (usually the Resource Centres) to the Accounting Portal and other aggregation data endpoints (see below). This includes support for:

- Accounting for Cloud computing in EGI (see also section 2.2.2.3)
- Storage consumption accounting
- Accounting MPI parallel job usage
- Application usage
- Transmitting pre-aggregated Accounting data instead of raw records.

Per-User aggregation: Supporting EGI's general direction of improvement of user experience, user-centric accounting data aggregation becomes necessary. This will most likely include new aggregation algorithms in the central (and regional) accounting repositories, as well as new views on the accounting central and regional accounting portals.

Support for billing: In collaboration with the Pay-per-use task force the accounting data will be assessed for its fitness for use in a billing subsystem.

2.1.1.5 Messaging

The EGI Messaging infrastructure is a legacy component of SAM, which is used to provide the Monitoring service in EGI (see section 2.1.1.3): Although it is no longer exclusively used by SAM, the ownership and development effort still lies within the Monitoring service development team. Currently, it is used by the Monitoring, Accounting and Operations Portal subsystems. The following activities are planned for PY4:

Secured Messaging connections: With the expansion in the use by more than one component, access control to the messaging network need to be secured. This will include transitioning clients to authenticated connections, tracing unauthenticated hosts, as well as SAM publishing authenticated test results to any number of authenticated consumers.

Messaging endpoint failover: Particularly targeted at the Messaging service, failover-enabled Messaging endpoints allow a consumer of this service to still use the Messaging infrastructure (via an alternative endpoint) if the initially requested endpoint is interrupted or non-reachable.

Monitoring the messaging infrastructure: The messaging infrastructure needs to be monitored to ensure satisfactory service delivery. To avoid circular dependencies with SAM, GANGLIA will be considered as a monitoring tool for the production messaging infrastructure.

2.1.1.6 Central Service Registry

EGI's central service registry is using the Grid Configuration Database (GOCDB) to store topology and service distribution data across its production infrastructure. Information stored in GOCDB is increasingly accessed by other EGI services such as Monitoring, Service Desk, Operations Portal, etc. Similar to many other services deployed as part of any of the EGI platforms, GOCDB seeks to improve sustainability through a number of measures as outlined below:

Major release GOCDB v5: The new major release of GOCDB actually introduces a number of new features and new technologies that are all required to implement a number of further improvements. Therefore, GOCDB v5 bundles a number of changes, among which the following are key:

- Replace legacy PROM DB with an O-R Mapping layer, making GOCDB independent of the persistence layer
- Simplify DB logic and queries; this reduces maintenance cost, simplifies future developments and last but not least improves query performance.

Extended Scoping: Currently, GOCDB only allows exclusive “EGI” and “Local” scoping of any DB entries. This limitation will be lifted in favour of non-exclusive, flexible tagging logic. This will reduce and prevent data duplication since a data/site entry can be tagged multiple times (e.g. for multiple projects).

Extensibility / GOCDB as a Service: Later updates to GOCDB 5 will allow for custom entity types in the form of key-value pairs. Together with the extended scoping feature described above, this will allow GOCDB host multiple projects and customers, so that GOCDB can serve as solution underpinning a Service Registry service to an extended customer base.

GLUE2 Information Model support: GOCDB will support the GLUE2 information model in the future. While the internal information model and rendering language still remains Doctrine, GOCDB’s external interface and integration will allow rendering queries in GLUE2 XML, as well as feeding GOCDB with updates expressed in GLUE2 XML rendering.

GOCDB failover support: The GOCDB failover configuration will be improved in order to allow automated switching to the back instance in case a failure is detected affecting the primary instance.

2.1.2 Support Services & Tools

2.1.2.1 Operational security

Operational security is an activity that is particularly crosscutting in nature and many detailed activities contribute to the overall effort matching the challenge of securing a highly distributed and federated production infrastructure. In general, most of the day-to-day work of operational security is well defined and will continue to be followed following the well-established Security Officer on Duty rota³. On top of these, the operational security activities for PY4 are grouped as follows:

Security monitoring tools: Security monitoring will be a prevalent theme in PY4 for operational security services, contributing mostly to automation and improved overseeing and follow-up of security incidents and vulnerabilities. These include:

- Requirements will be gathered requiring the extension of the security dashboard provided by the operations portal, in order to assist NGIs in the migration to SHA-2 compliant software and middleware upgrades.
- A proposal for site-wide monitoring will be presented: The document will define a plan for OMB discussion and the running of a technical pilot with a few NGIs. Results of the pilot will be implemented and evaluated.

³ https://wiki.egi.eu/csirt/index.php/Security_Officer_on_Duty_tasks



- A number of software patch status monitoring services will be deployed: An EGI instance will be primarily supported focusing mainly on sufficient support for site-wide monitoring.
- Probes: the SVG/CSIRT issues/alerts will be assessed to make sure they're monitored. Sending of notifications will be enabled.
- Collaboration with dashboard developers: Operators will be involved in the handling of non-critical issues (in order to do so, non-critical results have to be displayed in the operational dashboard).
- Reporting: Reports should be generated for discussion with the project management, NGIs and Resources Centres. Reports will provide periodic plots summarizing number of issues detected/handled.
- Monitoring: Central monitoring of security tools needs to be implemented.
- Nagios: CRL checking on services that have gridftp (CEs/SEs) and checking for known vulnerable file permissions via gridftp.

Incident response (IR) management collaboration: Collaboration with EUDAT and PRACE will continue and intensify to establish a cross-infrastructure operational security in response to user communities increasingly spanning multiple e-Infrastructures.

Improved IR ticketing support: Several improvements to the underlying ticketing system, including EGI customisations will be deployed; a change of authentication infrastructure integration away from EGI SSO might be necessary.

EGI and NGI security drills: The results of SSC6 will be evaluated and used to design, implement and execute SSC 7 in PY4. NGI-level SSCs will continue as required and requested by the NGIs.

EGI-wide support and enforcement for SHA-2: Affecting all aspects of the production infrastructure, services and subsystems need to be upgraded to transition from SHA-1 (vulnerable to attacks) to SHA-2 by the end of September 2013. This includes monitoring services that are *not* supporting SHA-2, and providing support and consultancy to user communities.

Review and update of security procedures: Two security procedures will be finalised (central compromised user handling) and defined (compromised certificate handling). In this context, a general review of the incident handling procedures will be conducted.

Training & dissemination: Further security training for NGIs, Resource Providers, and User Communities will be provided at the EGI Technical Forum 2013.

2.1.2.2 Operations Portal

The Operations Portal has served the EGI community for many years, and with the forthcoming changes in the EGI community at large a number of changes are necessary that all together will improve the sustainability of the tool itself as a pivotal means of overview of the EGI operational community.

Dashboard refactoring & framework upgrade: The dashboards will undergo a refactoring by transitioning to using Twitter's Bootstrap CSS framework for the various dashboards, allowing implementing the following new features:

- New, uniform look and feel across all dashboards

- Improvements in efficiency, reactivity and visual appeal of the dashboards
- A complete history is available for an issue
- Automatic removal of alarms with ok status in the main view
- Alarms grouping mechanism with possibility to add /remove alarms after the ticket creation
- Detect issues on local or non-EGI resources

Service desk support through plug-in architecture: Support for multiple service desk solutions will be implemented using a plug-in architecture combined with configuration files mapping solution idiosyncrasies into the generic portal model. The default service desk integration will be GGUS, used for EGI's service desk service.

VO-oriented A/R reporting: Improving the user experience, availability and reliability (A/R) reports will be generated from a VO point of view. Serving so-called "high-activity" VOs, the generated reports will regard only those resources the respective VO has access to.

Continuous integration & modularisation: The Operations Portal will be further modularised leading to a clearer architecture and overall maintainability. Regular updating to latest Web technology and social interaction services (e.g. RSS, widgets) completes the overhaul of the Operations Portal.

2.1.2.3 Metrics portal

While not meant to be consumed by EGI research communities and platform integrators, it develops into an indispensable tool for federated e-Infrastructure operations management. This is reflected in the planned improvements:

Look & Feel Improvements: The UI needs to be improved to accommodate its more central role for EGI operations. This includes improvements in the presentation layer for editing, reporting and charting functions, as well as corporate identity integration.

Mobile access: Mobile access is a frequently needed functionality, particularly when figures for certain KPIs are needed on-demand in situations where laptops are inconvenient or not an option.

Extended EGI Service Desk metrics: The EGI Service Desk plays a pivotal role in EG, both for operational purposes and as a service to user communities. As such the Metrics portal will provide extended support for KPIs pertinent to the EGI Service Desk.

Better support for manual metrics: Manual metrics are KPIs that are entered manually by a user who is responsible for collecting the respective metrics. Those metrics are often difficult or impossible to collect in an automated way, so the Metrics Portal must accommodate maintaining such metrics.

Customised reports & Excel export: Often reports on metrics are needed in similar, yet different aspects. Customised reports allow the user to tweak certain parameters of reports to achieve this without having to implement each report variation. Also, reports including customised reports will be available for Export into an Excel spread sheet format.

2.1.3 Software Provisioning

In EGI, Software being part of any Community Platform that is to be deployed directly onto the physical infrastructure must undergo a software provisioning process. Currently, this process is



applied by EGI.eu and its partners to all software that is published as part of the Unified Middleware Distribution (UMD) in the EGI Software Repository. This process may need to be extended to other community platforms. This provisioning process is modelled after common Software Acceptance Testing methodologies and encompasses three major services: (1) Quality Assurance, (2) Quality Control, (3) Staged Rollout, and using the integrated Software Repository (see section 2.3.1.3) as the major integration point with the Resource Centres who roll out the software into the production infrastructure.

Over the last three years the Software Provisioning service mainly focused on Quality Criteria pertinent to the software coming from the IGE and EMI projects that served as the principal Technology Providers to EGI. With the conclusion of those projects, and the general move towards a platform-oriented architecture for EGI's production infrastructure, the provisioning activities will adapt to the new situation as follows:

2.1.3.1 Quality Assurance

Adapted Quality Criteria review process: Representatives of the IGE and EMI projects used to review the Quality Criteria and provide their feedback before the new editions were finalised and put into force. With the end of these projects, new reviewers and an adapted review process will accommodate this situation.

Quality Criteria prioritisation: Not every Quality Criterion needs to be checked for every software release that is taken through the provisioning process. Criteria will be arranged into mandatory and optional testing groups for major, minor and revision releases of the pertinent software.

Quality Criteria compendium: The EGI Wiki will be used to build a collaborative Quality Criteria compendium next to the formal testing documents.

Test automation: Where possible and feasible, canonical, reviewed, and automated tests will be provided as part of the Quality Criteria documentation, helping both Technology Providers and Quality Control Teams to achieve consistent high-level test coverage for provisioned software.

Two new revisions: As per release schedule, two new revisions (no. 6 and 7) of the Quality Criteria documents will be released to Technology providers and Quality Control teams within EGI.

2.1.3.2 Quality Control

Quality Control requires a test infrastructure where easy and trace-less service deployment with minimal effort is essential – Cloud Computing is a principal enabling technology for these goals. In this sense the Quality Control team will integrate with the EGI Cloud Infrastructure platform as follows:

Consistent test bed management through OCCI: The Quality Control team has access to a private institutional Cloud for deploying their test infrastructure. Deploying and integrating with the rOCCI toolkit not only allows scaling out and automating the test infrastructure management, but also prepares this private Cloud to be integrated into the EGI Cloud Infrastructure Platform.

VM Image management automation: By integrating with the vmcaster/vmcatcher technology being used in the EGI Federated Clouds Task Force, a common VM image management infrastructure will be available which allows the Quality Control teams to re-use stock VM images necessary to test an



updated software component. This includes easy and streamlined distribution of updates to the stock VMs that are necessary over the course of time.

Test-bed provisioning automation: Often, software that is verified against applicable Quality Criteria does not require a full clean install, but an update or a change in its configuration from a previously successfully verified version. By exploiting VM Contextualisation proven in the EGI Federated Clouds Task Force, provisioning of a test bed will be a suitable target for automation, further increasing efficiency of the overall software provisioning process.

2.1.3.3 Staged Rollout

The planned changes in the architecture of the EGI production infrastructure, and the end of the EMI and IGE projects require a number of changes in the Staged Rollout part of the Software Provisioning activity as follows:

Re-balance Staged Rollout effort distribution: Staged Rollout activity at the participating EGI Resource Providers is mostly contributed to EGI's Software Provisioning as community effort. This best effort needs to be re-balanced so ensure that the critical services are well tested.

Review Staged Rollout procedures: The end of the EMI and IGE projects coincide with the recent development of the EGI Cloud Infrastructure platform (see section 2.2). In conjunction with community effort distribution assessments, the general procedures for Staged Rollout need to be reviewed for possible benefits using the EGI Cloud Infrastructure platform for Staged Rollout tests.

Participate in EGI URT release scheduling: EGI.eu is taking on some of the release scheduling activities that were previously conducted by the EMI and IGE projects. Staged Rollout is an important coordination partner in this effort.

2.2 Cloud Infrastructure Platform

Deployed on top of the EGI Core Infrastructure Platform, the EGI Cloud Infrastructure Platform provides “*a federated IaaS Cloud infrastructure*” based on institutional private IaaS Clouds. It wholly embraces the Cloud paradigm and extends it with a federation mechanism that is partially based on the EGI Core Infrastructure Platform, and partially provides new federation and distribution services geared towards Cloud computing.

The Cloud Infrastructure Platform is a young addition to the EGI ecosystem, and needs time to mature towards integrating it into the EGI production infrastructure. Naturally, the overarching theme of the roadmap for the Cloud Infrastructure Platform is the preparation towards production inclusion in the coming year.

2.2.1 Federated Cloud management services

2.2.1.1 VM Management

In the EGI Cloud Infrastructure Platform, managing Virtual Machines (not the services and software that are deployed *within* that virtual machine) is done through the standardised OCCI interface. In principal, this will provide interoperability across different implementations and easier portability of integrations from one Cloud provider to another. Since not all Cloud Management Frameworks within EGI natively support OCCI, this gap is bridged with the rOCCI framework and rOCCI-Server, a



server-side implementation of OCCI and an existing integration with OpenNebula. To further improve the VM Management interface the following will be rolled out over the course of the next year:

Consistent Authentication using Grid technology: By integrating with the X.509 based PKI used in the existing EGI production infrastructure, a large number of EGI research communities will be able to access the EGI Cloud Infrastructure without having to obtain new credentials – and use the Cloud infrastructure the same way they use the EGI Grid infrastructure.

Re-design of rOCCI-Server: The current architecture of the rOCCI-Server has reached its limits in terms of scalability, maintainability and extensibility. To support integration with more diverse Cloud Management Frameworks, the rOCCI-Server needs a ground-up re-design and implementation.

OCCI client library in Java: To foster broader uptake of the OCCI interface on the client side, a Java library will expose an API that is capable of marshalling and un-marshalling OCCI messages (including the Grid Authentication extensions) in a conversation with an OCCI-based Cloud Management Framework.

2.2.1.2 Data Management

Standards-based Cloud Storage and Data management in EGI has a very similar architecture as the VM Management service: A proxy-server is fronting actual Cloud Management Frameworks that do not support CDMI (the chosen standard in EGI) as their access interface. Based on an existing prototype implementation, the following new features are foreseen in the coming year:

Consistent Authentication using Grid technology: To provide consistent Authentication mechanisms across all federated institutional Clouds, the CDMI proxy server needs to integrate with the current Federated AAI deployed in EGI.

CRUD operations for Object Storage: Based on the existing Cloud Object Storage feature, and the integration with EGI Federated AAI, full support for CRUD operations (Create, Read, Update, Delete) will be rolled out over the course of the next year.

Block storage support: Block storage is an increasingly popular means of Cloud Storage in that it provides storage as a virtualised device that fits perfectly with the device virtualisation paradigm used in Cloud Computing. Users will be able to manage block devices (e.g. create, delete, grow, shrink) and attaching available block devices to OCCI-compliant VMs via the CDMI interface. The exposed functionality will depend on the underlying Cloud Management Framework, with features outside the least common denominator being available through proprietary means only.

Extend CDMI metadata support: Particularly in Cloud storage scenarios, metadata plays an important role for not only miscellaneous information about the data itself but also for the automation of data management in larger usage scenarios. Extending the support for metadata that is mandated by the CDMI specification, and metadata that is necessary for CRUD operations on object storage operations (see above), additional metadata will be supported for a number of use cases. Candidates for metadata support (pending analysis) are:

- Geographic data placement
- Data retention policies
- Backup policies



JavaScript based CDMI storage browser: Users will get early and lightweight access to CDMI-based Cloud storage within and external to the EGI Cloud Infrastructure using a CDMI based Cloud storage browser.

SDKs in Python and Java: To facilitate 3rd party tool and platform integration with the EGI Cloud Infrastructure Platform, SDKs for Python and Java will be made available.

Accounting integration: Accounting for Cloud storage is necessary for it being integrated into the EGI Cloud Infrastructure Platform. This solution will integrate with the EGI Accounting service provided as part of the EGI Core Infrastructure Platform.

2.2.1.3 Image management and distribution

VM Images, once made available for use in the Cloud Infrastructure Platform, need to be distributed to the supporting Resource Providers who then in turn may conduct discrete tests before acceptance for instantiation. Also, this capability needs to support the distribution of VM Image updates so that older versions can be superseded. The following enhancements are foreseen:

Full deployment of vmcatcher/vmcaster in the testbed: The vmcatcher/vmcaster software toolkit was evaluated by a subset of Cloud resource providers, and it was accepted to be included into the regular toolchain of the Cloud Infrastructure Platform. It now needs to be deployed at all federated Cloud resource providers.

Integrate with the OCCI client libraries: Integrating with the OCCI client libraries allows for transparent selection of images and sites for managing the virtual infrastructure on top of the EGI Cloud services.

Integrate with the EGI Federated AAI: Only allow certain users in certain groups (i.e. VOs) to publish VM images and image updates.

Support single VM Image updates: Formally, each VM image has a unique identifier (a UUID). Consequently a second VM Image that constitutes an update of the first VM Image must have a different unique identifier, yet be flagged as a successor of the first VM Image, so that updates to a VM image can be distributed and managed in an automated way.

Support multi-VM image appliance updates: Frequently though not always one VM Image contains software bundled into one “Appliance”. More complex Appliances may require bundling its software into more than one distinct VM Image (e.g. a “head VM” and a “worker VM”). The VM Image distribution infrastructure consequently needs to support distributing updates to Appliances (both full updates and partial updates to the VM Image set).

Integration with AppDB: To further increase the attraction of the Cloud Infrastructure Platform, User generated scientific software may be registered in the AppDB as available in pre-packaged VM images that then can be distributed through AppDB acting as a VM Image publisher towards the Image Management infrastructure.



2.2.1.4 Information Discovery

Information Discovery in the EGI Cloud Infrastructure Platform is used to provide a mix of static and semi-static information about a resource provider in the context of Cloud computing. The purpose of this service is to describe the technical details of a Cloud resource provider's service offering, either as complementary information to the baseline federated information (for example, which Cloud Management Framework in which version is deployed), or as a means to describe specialised services that provide added value to a subset of the EGI Cloud infrastructure research communities (e.g. specialised AAI integrations, special VM image audit and endorsement procedures).

Effectively re-using components of the EGI Core Infrastructure Information Discovery subsystem the counterpart in the EGI Cloud Infrastructure Platform is using a scaled-down version of the full Information Discovery subsystem deployment. The goal is however, that over time this can be replaced by integrating with the existing Information Discovery subsystem coming from the EGI Core Infrastructure Platform. The reason for this parallel deployment is at the same time the topic for future plans in this area:

Finalise Cloud-related GLUE2 extension: Related and synchronised to the EGI GLUE2 profile (see section 2.1.1.2) initiative, extensions to the GLUE2 information model that properly describe Cloud-specific resources need to be finalised, integrated into the profile and eventually implemented.

2.2.2 Core Infrastructure integration

As described earlier⁴, each higher-level platform (be it a Community Platform or the Clouds Infrastructure Platform) must integrate with a number of capabilities and services provided by the EGI Core Infrastructure Platform. The EGI Cloud Infrastructure Platform plays a pioneering role in this context, as it is the first well-defined platform within the EGI community that integrates in the desired way with the EGI Core Infrastructure Platform.

Extend site certification procedures with Cloud aspects: EGI formally certifies sites ready for integration into the federated production infrastructure. This procedure needs to be extended to be able to include sites into the production infrastructure that expose Cloud resources. Volunteer Cloud Resource providing sites will work with the EGI Operations team to evaluate the necessary changes, work with the OMB to approve these changes, and eventually start taking sites through the certification process into production.

2.2.2.1 Virtual Organisation management & AAI

Virtual Organisation management is a required feature in a federated, distributed infrastructure that is used / shared between multiple user communities, and this is no different for a federated Cloud infrastructure. As this is a cross cutting activity, the integration work with the other management interfaces and capabilities are not mentioned here.

Uniform X.509, RFC Proxy certificate and proxy-of-proxy certificate support: Principle X.509 certificate based user authentication and authorisation works in a federated environment, but "falls apart" at the fringes and certain corner cases that are frequent in the Grid infrastructure, for example using proxy certificates that were issued using a proxy certificate. Also, full RFC proxy style authentication is not fully available in the whole federated Clouds infrastructure. Enhancements and

⁴ More detail on the mechanics can be found in MS514: EGI Platform Roadmap (see section 6)



changes to existing Cloud Management frameworks will be deployed over the next year to address these issues.

Multi-VO support: Before being used in production, support for more than one VO across the different Cloud Management Frameworks must be implemented uniformly; particularly VO membership across Cloud providers must be maintained so that resource usage per VO will be correctly accounted for.

2.2.2.2 Monitoring

Currently the integration with the Core Infrastructure Platform Monitoring Service is limited to operating an independent, production quality SAM instance (see section 2.1.1.3 for more information). A number of custom, Cloud-specific Nagios probes are deployed monitoring the general service reachability, OCCI availability and reliability, EGI Accounting compliance, and Information Discovery service compliance. The following enhancements are planned to prepare the Cloud Infrastructure monitoring for regular production use:

Implement and deploy a Cloud Storage probe based on CDMI: As EGI is extending its federated Cloud infrastructure to support Cloud Storage (see section 2.2.1.2) these Cloud storage services need to be monitored as well. A monitoring probe will be developed and deployed in the current independent test-bed SAM instance.

Integrate Cloud probes in production Monitoring releases: As part of the site certification procedure all or a subset of the current Cloud monitoring probes need to be integrated into the production SAM releases that are deployed in the NGIs; also, the monitoring profiles for NGIs need to be updated so that NGIs can enable or disable Cloud monitoring when providing production Cloud infrastructure within EGI.

2.2.2.3 Accounting

Accounting for Cloud resources is currently at an early stage. Based on a small subset of Grid Compute accounting (i.e. Number of cores, wallclock time, allocated RAM, storage for the VM image), the infrastructure has been set up to collect Cloud accounting data from the various Cloud Management Frameworks, and to persist and aggregate them using the current production Accounting service (see also section 2.1.1.4). Nonetheless, until the current solution has been certified for production use, Cloud accounting data is processed using the same infrastructure, but separate persistence entities so as to not interfere with production-grade accounting. The following is anticipated for preparing the Accounting integration into a pre-production state:

Extend Cloud accounting to cover Cloud Storage: It is necessary to account for the consumption of Cloud storage just as much for Cloud Computing. An initial set will include both temporal and spatial aspects, i.e. differential accounting in a similar model to the calculation of annual interest rates.

2.2.2.4 Central Service Catalogue

From the viewpoint of the EGI Core Infrastructure, the EGI Cloud Infrastructure is just another Community Platform and thus required to register the Cloud resources in EGI's Central Service Catalogue. Since all necessary service types are already available in the current service, no further integration work with the Central Service Registry is planned.



2.3 Collaboration Platform

The EGI Collaboration Platform includes services that are shared by all EGI users, and some of the other EGI stakeholders. The services are common to different communities, but are not critical to the operation or use of the EGI production infrastructure. It thus complements the EGI Core Infrastructure and EGI Cloud Infrastructure platforms, and contributes to their efficient use.

The services included in the EGI Collaboration Platform fall into two principal categories as follows:

- **Technical collaboration services** connected to components of the EGI production infrastructure (i.e. the EGI Core Infrastructure, EGI Cloud Infrastructure or any other Community platforms)
- **Social collaboration services** supporting information and knowledge exchange within EGI, and between EGI and members of the ERA.

2.3.1 Technical collaboration services

2.3.1.1 Service Desk

The EGI Service Desk is a service offered based on the Global Grid User Support (GGUS) system. Reflecting the on-going reorganisation of the EGI production infrastructure into a platform oriented approach GGUS will further evolve into a more complete Service Desk solution by implementing the following:

On-demand reporting: The legacy report generator will be replaced with a completely rewritten alternative that is more aligned with operational needs; the work already begun in PY3 and will be finalised in PY4.

High-Availability for GGUS: The service desk is pivotal for an excellent user experience and therefore any service outage will immediately diminish user satisfaction. High-availability will be achieved through integrating Virtualisation technology into GGUS.

Disaster management: Complementing the High-Availability activities, disaster management will be put in place as an additional risk management activity.

Shibboleth integration: Many users of GGUS have requested alternative authentication methods to the current username/password and X.509 based methods. Shibboleth has been identified as the solution with the broadest support of the existing customer base.

Alarm ticket improvements: Alarm tickets bypass the normal service desk workflow, and allow notification of support teams irrespective of office hours. Access to this feature is naturally very limited; RBAC (role based access control) will be improved through fail-safe integration with VOMS. Also, alarm tickets will be enabled and integrated for EGI Operations staff.

Revision of software support processes and technology helpdesk: GGUS will implement a new notification mechanism that periodically reminds third level software support staff in case feedback is requested from them. This is needed in preparation to a scenario with decreased support effort, where software support is no longer coordinated and supervised at a project level and is completely delegated to support teams. All software support units will be re-assessed in terms of quality that can be guaranteed in PY4. Two different service quality targets will be proposed. The access rights to technological helpdesk will be revised to facilitate cooperation with teams. GGUS support will be



integrated with other external helpdesks/support channels for those products whose support teams are not directly accessible in GGUS.

2.3.1.2 VM Marketplace & Appliance Repository

The VM Marketplace is a metadata registry for publically available VM images while the binary data of the images themselves are stored in Appliance Repositories elsewhere. Thus the central registry service is backed by distributed VM Image storage services, from which the images are distributed to the participating Cloud Resource Providers.

For the time being, there are no improvements foreseen for the VM Marketplace.

2.3.1.3 Software Repository

The EGI Software Repository is designed as the principal source of software to be deployed in the EGI production infrastructure, next to the base Operating System repositories. As such, it is hosting the Unified Middleware Distribution (UMD), which in turn is an amalgamation of currently predominant Grid Middleware (i.e. Globus, Unicore, dCache, ARC, and gLite) deployed in the production infrastructure. With the end of the IGE and EMI projects, the EGI Software Repository will continue to host the UMD and other EGI-specific software. The following changes will prepare the EGI Software Repository for the time beyond the EGI-InSPIRE era:

On-demand repository services for community software: Integrating with the Application Database (see section 2.3.1.4) the EGI Software Repository will become the host for software that is developed in the community, for the community.

Provisioning-as-a-Service: With EMI and IGE projects coming to an end, EGI is preparing for changes in the landscape of its Technology Providers. It is clear that not all software coming from all affiliated Technology providers can be passed through the Software Provisioning process used for the Community Platform integration service (see section 2.1.3), nor is this desired. Software that is neither community-based (i.e. provided through the Application Database) nor passed through the Software Provisioning Process, a third, lightweight software provisioning process will be provided.

2.3.1.4 Application Database

The EGI Applications Database (AppDB) is a centralised service that stores and provides information about a rich variety of software that has been integrated with one or more components in the EGI production infrastructure. It facilitates re-use of scientific software and associated tools by providing a platform for developers to publish new and updated versions of their software, integrated in the backend with the EGI Software Repository for easy access to packages. As with many other tools and services in EGI, the AppDB service in the next year will receive much attention for its sustainability options and future use:

Support for VM images: Extend the AppDB infrastructure so that developers can publish “Appliances”, i.e. ready-made Virtual Machine images (VMIs) that can be downloaded and deployed without much further effort.

Integrate with the EGI Cloud Infrastructure Platform: Taking VM Image publishing one step further, integrating AppDB with the EGI Cloud Infrastructure Platform’s VM Image distribution infrastructure is all but a logical next step. This greatly contributes to the automation of EGI’s infrastructure management and improves the overall user experience.



AppDB as a service: Further maturing AppDB into a service offering beyond EGI to research e-Infrastructures of similar size allows diversification of the customer base securing steady sources of income and lowering the team's dependence on single-source funding from the EGI-InSPIRE project.

2.3.2 Social collaboration services

Out of the currently eight social collaboration services only two are included in this roadmap, as described below. The other services are already in a very mature state and will be maintained as required on the current level.

2.3.2.1 Training Marketplace

The EGI Training Marketplace (TMP) is an online registry to advertise and learn about (browse and search) training on software and services within EGI. This includes online training material, and events such as university courses. During PY4 the Training Marketplace will mainly focus on sustainability through broadening the customer base by promoting the use of the Training Marketplace web gadget to scientific user communities within and beyond EGI.

The following developments are planned to support this goal:

TMP as a service: The TMP software tool will be turned into a service, enabling the provider to expand the customer base. By deploying a UK-specific training marketplace based on the EGI Training Marketplace STFC will learn the specific requirements and necessary changes to turn the tool into a service.

Web gadget theming: Customise the web appearance of the web gadget to better fit into external websites.

New content type: Add “webinar” as content type reflecting the increase of this type of training in the community

Free tagging & tag filter: Allow any type of tag on Training Marketplace entries, with the web gadget to filter on specific tags for display.

2.3.2.2 Client Relationship Management

The Client Relationship Management (CRM) service has been setup to help EGI coordinate and monitor its outreach activities targeting new user communities. Individuals from the EGI community record key findings about the e-infrastructure needs of potential clients of EGI services. The EGI CRM targets the NGI International Liaisons, and the EGI Champions, as well as the EGI.eu marketing and user community support teams.

Data hygiene: Increase the accuracy and completeness of information stored in the CRM, focusing on ESFRI projects' interesting e-Infrastructures and potential for them engaging with EGI.

Maintenance: Based on the Open Source vTiger software, the CRM will track and include new developments of vTiger and any new requirements regarding collecting useful information on ESFRI user communities.

ToDo lists & reminders: Grouped around the information gathering activities, basic ToDo lists and monthly reminder Emails will be added.



Improved UI: Contacts information for individuals and organisations will be improved for a better user experience.

3 COMMUNITY & COORDINATION

The second pillar of the EGI Strategy, Community and Coordination (see Figure 3) complements the Operational Infrastructure and the VRE support spanning all levels and aspects of the EGI community. From reaching out to system administrators that manage the physical infrastructure and deploy software components coming from all types of software platforms to connecting researchers with colleagues in the field to examine the feasibility and synergies in establishing shared Virtual Research Infrastructures, community and coordination activities provide the cohesive human network required to nourish the EGI community at large.

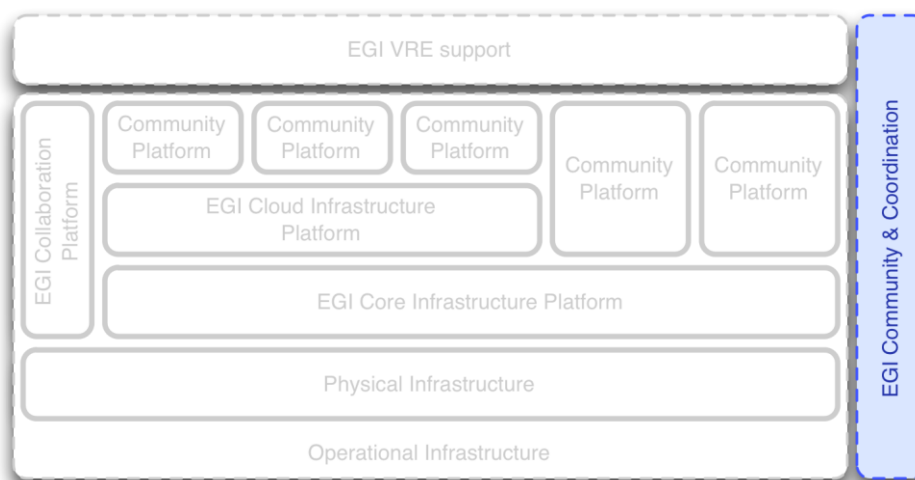


Figure 3: EGI Community and Coordination services form the second pillar of the EGI strategy.

EGI Community and Coordination activities are organised in three areas as follows:

Human networks – Similar to the human circulatory system the EGI human networks permeate the EGI community on all levels, and through its activities connects and establishes communication paths between whoever needs to talk to anybody else: Connecting people at its best.

Communications and marketing – While human networks tend to orientate more towards the inner workings of the EGI community, communications and marketing activities provides external facing services: press contacts, press releases, conventions, event planning, promotion material and social media interaction.

Policy and Strategy – The third component of EGI’s community and coordination activities provides the formal framework and direction towards reaching EGI’s vision for H2020.

3.1 Human Networks

In 2011 the EGI-InSPIRE project revised its user engagement activities and defined a new structure for the NA2 work package within an updated Description of Work. The purpose of the restructuring was to improve the efficiency and flexibility of the interaction between the NGIs, EGI.eu and other organisations to achieve common goals in the field of "Engaging with New User Communities". The new structure consists of two key elements: NGI International Liaisons and a Virtual Team framework. Both of these initiatives proved to be highly successful in that they transformed the relationship between EGI and its users from a “Supplier says to Customer” to one to a more engaging “Supplier *listens* to Customer” relationship. Taking this relationship even one step further, EGI and its



customer now have a framework within which to firstly ‘listen’, then to collaborate and let the customer lead in developing solutions for the community at large.

The success of the past year’s work by NILs and VTs has largely been attributable to the formation of human networks across the grid community, people who are able to communicate and interact effectively and collaboratively for the common good. Recognition of such positive progress has led to further initiatives that build on human networks: firstly the EGI Champions scheme which was kicked off in September 2012 during the Technical Forum in Prague. Secondly, an evolution of the VT project scheme christened as EGI Mini-projects was launched during the Community Forum in Manchester; Mini-projects are run by the community for the community and are underpinned with EGI-InSPIRE funding, thus increasing the scope and size of projects that can be run. Further human network initiatives will be launched in coming months:

- A ‘Geek Squad’ network will be established that nurtures the technical expertise residing in the product teams that provides solutions to EGI’s production infrastructure and its user community. This could include support for travel to EGI events and to engage in gathering detailed requirements for new functionality from research communities.
- An Operations network that brings together NGI and resource centre operations staff will be established with the objective of forming highly competent groups of technical experts who can help resolve problems for ordinary users, either in terms of specific computing issues or in terms of service delivery and service management issues.

3.1.1 NGI International Liaisons

The interaction between the NGI teams and EGI.eu on non-operational activities is undertaken through a network of "NGI International Liaisons" (NILs), nominated individuals who are responsible within the NGIs for the delivery and interaction of non-operational tasks. Non-operational activities cover areas such as marketing & communication, strategic planning and policy support, community outreach and events for new users, technical outreach and support to new communities. The role of the NIL recognises both the complexity and diversity of each NGI and also the need for these NGIs to be involved in the pan-European, coordinated, non-technical activities of EGI. It is not necessarily the NILs who undertake any of the non-technical activities but instead, they make sure the appropriate individuals or teams within the NGIs respond to any particular activity or issue that arises.

Improvements continue to be made in the support given to the NILs and other networks, including:

- Monitoring and review of the effort recorded by NILs through their timesheets and their impact.
- Continued clarification and sharing of best practice across the NILs with each other.
- Exploring the use of social media to help develop and strengthen interaction between the NILs and other networks.
- Introduction of monthly telecons and webinars on topics relevant to the NILs and the communities they represent.

3.1.2 EGI Champions

The broad concept of establishing a cadre of ‘ambassadors’ who would act as enthusiastic and proactive promoters of EGI was launched under the banner of “EGI Champions” during the EGI Technical Forum in Prague in September 2012. Since then, the practical implementation of this concept has been set in place, firstly via the [EGI Web site](#) and [Wiki](#) pages which provided greater detail to the new initiative and subsequently through the recruitment and selection of the 1st cohort of 6 EGI Champions in November 2012 followed by a 2nd cohort of 3 in February 2013.

The support given to the EGI Champions will evolve following the initial F2F meeting that took place at the EGI Community Forum in Manchester in April 2013. It is expected to include:



- Regular monthly telecons to enable bidirectional discussion between the EGI Champions and EGI.eu reporting on recent activities and developments within their country/community.
- Webinars relating to issues seen as a priority to the EGI Champions, such as training to improve communication skills and updates of the EGI and Community platforms.

The teleconferences will provide opportunity for the NILs and for the EGI Champions to.

3.1.3 Operations network

Operations coordination drives through the Operations Management Board future developments in the operations area by making sure that operations evolve with the needs of the community and to support the integration of new resources and middleware platforms (e.g. desktop grids, virtual machines, high performance computing). It does this by providing coordination and management and by developing policies and procedures for the operational services that are integrated into the production infrastructure through the operational support of distributed operations teams. Coordination of software deployment and feedback gathering is delivered through fortnightly operations meetings. In PY4 OMB activities will continue to run stably. Stronger cooperation with outreach teams will be sought for to ensure that successful testing activities for the support of new communities can be quickly integrated into a production environment.

The User Community Board will be re-scoped by giving more emphasis on the support of existing user communities and VRCs, by ensuring stronger collaboration with NGI operations.

OMB and UCB meetings will still be complemented by bi-weekly technical meetings dealing with software deployment issues, the output of staged rollout activities and UMD release activities.

3.1.4 Geek Squad

With the end of the EMI and IGE projects the support provided for interaction with research communities and the development of new technology has been significantly reduced. However, the human network and capital that currently exists within the product teams can still be supported and developed, and opportunities provided for their engagement with EGI.

The EGI Geek Squad (a concept that is still being developed) will build upon the experiences from the NILs and EGI Champions network to develop a loose collection of informally committed (by interest) group of experts coming from the technology community within the EGI ecosystem. It is envisaged that travel support could be provided to EGI meetings, engagement with research communities to gather new requirements, and their engagement in virtual teams and funded mini-projects (see below) to support limited technical developments and integration.

3.1.5 Virtual Teams & mini Projects

The human networks (NILs, EGI Champions, Operations and Geek Squad) represent a critical source of expertise and experience that EGI needs to bring together to tackle important community issues through unfunded Virtual Teams and funded Virtual Teams (mini-projects). The impressive results of the VT project scheme depend on the positive and proactive approach of all those involved in taking an embryonic requirement or idea through to the delivery of real and practical results in under 12 months.

Virtual Teams (unfunded): Nearly 15 VT projects were undertaken and completed in 2012. At the time of preparing this report, a further 4 have been completed or are in course of closure and 4 more



are under way. The status of the projects is presented via the EGI VT Wiki⁵ pages and weekly progress reporting is compiled and provided to the EGI.eu Director and senior staff.

The VT framework has proved to be a very effective method for undertaking small projects for the benefit of the EGI ecosystem but the effort expended has been unfunded. This means that to a large extent, VT projects have been run on a ‘best effort’ basis. In most cases this has imposed significant burdens on some participants while others participants have been unable to contribute as much as they had initially intended. EGI.eu’s TONC team has already seen steady improvement in the structuring and management of new VTs, mainly as a result of initiatives to provide project management consultancy and assistance through the project start-up phase and in planning and conduct of project progress meeting. This approach will continue to be developed through PY4 with Lessons Learned being fed back into the VT process – template project documents, procedures, assistance in the use of available tools (e.g. Webex teleconferencing) and advice. Tighter management with better assistance from the very start of projects right through to their conclusion will help ensure VTs are more able to focus on what needs to be done and by who.

Virtual Teams (funded – a.k.a mini-projects): For PY4 a new VT model ‘funded Mini-projects’ are being started from April 2013. At this stage there are 11 approved ‘funded Mini-projects’ which are already being formed using the VT framework – the EGI “Funded Virtual Team project” Wiki pages⁶ have been established and guidance on reporting procedures and available support was issued during a dedicated session at the EGI Community Forum 2013 in Manchester. Funding to the extent of over €640K has already been identified for these 11 Mini-projects through the EGI-InSPIRE project.

3.2 Communications and marketing

EGI needs to “go to the user” and this requires engaging disciplines with targeted content and information provided through specific channels. This will mostly be through event attendance and material production. The communications team will also engage the other audiences through similar channels and products.

3.2.1 Major EGI Events & Workshops

The communications team will support the two annual EGI Forums and other internal events. Its main focus will be on supporting EGI involvement in events aimed at individual scientific disciplines. The team will work with partner projects, VRCs and the NILs and Champions networks to increase EGI’s visibility at national or international events. This could be through helping with content for a Champion’s presentation at a local seminar to funding a booth at a major at international event staffed by members of a VRC or a NIL. This will reduce travel costs for the central team and increase the number and range of events that can be attended.

3.2.2 Materials

The catalogue of materials produced in the last three years will be extended and refreshed as required. Both creation and maintenance of the products will be done with input from members of the EGI.eu staff and other stakeholders. In particular the Champions will be asked to advise on the content of material aimed at their community. Types of materials will include the “Stories from the grid” series, a portfolio of audience specific materials. It includes posters, brochures and videos aimed at the non-

⁵ https://wiki.egi.eu/wiki/Virtual_Team_Projects

⁶ https://wiki.egi.eu/wiki/Overview_of_Funded_Virtual_Team_projects



technical user⁷ and brochures aimed at policy makers such as ‘The Case for EGI’, based on the EGI Strategic Plan and the annual reports.

3.2.3 Main website and web design

The EGI website was re-launched in March 2012, to be more visually appealing and has a layout that is designed to encourage audiences to stay on the site to read the material in more depth. The team will continue to tweak the design and content to ensure visibility of the most important aspects of the site. The site also includes the ability for audiences to engage with both EGI and the content of the site through various social media features incorporated throughout the site allowing content to be shared easily. Moving forward the team will maintain the relevance and usefulness of the site by ensuring that more resources like case studies, images and guides to the infrastructure are added.

3.2.4 Publications

The EGI.eu communications team produces two regular internal publications, contributes to external publications and provides deliverables for the project. The Communications Team produces a quarterly newsletter aimed at the community, containing relevant information and links in an easy to digest format. The articles are available individually as pages on the website and promoted through the social media channels. Each month, the team produces a monthly Director’s Letter on behalf of the Project Director, which is issued to the all-project mailing list and published on the website.

3.2.5 Dispatches

The communications team produces a very short monthly email highlighting tasks, requests and information that is sent to the NILs. This “dispatch” was created as a result of feedback from the NILs and has proved very popular. The Communications Team are investigating creating similar dispatches aimed at the Champions (bi-monthly) and the developer community that use the recently launched Community Software Repository⁸ (quarterly).

3.2.6 External Publications

The communication team also targets articles about EGI and the scientific work achieved using the infrastructure at other trade publications and project newsletters. In particular, articles will be produced for *iSGTW*⁹, and to policy focused publications such as *The Parliament*, *Pan European Networks* and *Public Service Review*.

3.2.7 Media and public relations

EGI has renewed media partnerships with Tabor Communications for PY4, which publishes HPCwire¹⁰, HPCinthecloud¹¹ and Datanami¹². These publications publicise EGI events through web banners on their websites, and are invited to send journalists to report onsite from the events, which can lead to stories being picked up by the wider media. These partnerships will continue during Year 4, covering the EGI Technical and Community Forums or other relevant workshops. EGI will also

⁷ <http://www.egi.eu/news-and-media/videos/>

⁸ http://www.egi.eu/blog/2013/04/15/egi_applications_database_as_a_community_software_repository.html

⁹ www.isgtw.org

¹⁰ <http://www.hpcwire.com/>

¹¹ <http://www.hpcinthecloud.com/>

¹² <http://www.datanami.com/>



work to establish connections to individual journalists, for example through contacts established with journalists at the BBC World Service and the Times Higher Educational Supplement, which have been developed through attending and organising media training.

3.2.8 Social Media

In the last year EGI has been working on leveraging these newer channels for engaging both the existing community but also people outside its traditional audience. Although focusing on Facebook and Twitter, and its self-hosted blog, EGI also has a presence on Google+, LinkedIn, YouTube and Flickr¹³. In Year 4, the EGI social media presence will be strengthened and improved, particularly in the Policy, User Community, Technology & Operations areas which currently have lower engagement levels than the main social media channels. Working with the relevant teams at EGI.eu these will be used to form a basis for the community building efforts around EGI events, and will be used to launch competitions and viral campaigns.

3.3 Strategy and Policy

As the project moves into its final year, sustainability plans have become increasingly coupled with EGI's long-term strategy to connect researchers from all fields of science across the whole ERA with the reliable and innovative ICT services they need to undertake their collaborative world-class and world-inclusive research. Over the course of the next year, it will be essential to refine these strategies and long-term solutions to ensure EGI is sustained for years to come. This will be done through targeted strategic analysis, improving service management, developing relationships and business opportunities as well as policies for long-term sustainability.

3.3.1 Strategic Analysis

The wider e-Infrastructure community, whether directly or indirectly involved with EGI, benefits from informed decisions being taken. A range of actors relies on insight and data to support both strategic and tactical decision-making. This includes EGI management trying to make informed decisions, a site administrator looking to understand the impact of developments on the infrastructure, technology providers looking for new opportunities, and project managers reviewing product roadmaps. Specific activities will focus on:

Support the evolution of EGI's 2020 Strategy: EGI.eu Strategy and Policy Team analyses strategic themes and trends globally and in Europe and produce documents and reports to inform the EGI management bodies and wider community to support the decision-making process and articulate an evolving landscape for the benefit of the wider community. The SPT will continue to provide the EGI Council and Management bodies with policy papers to enable informed decision-making.

Update the EGI Balanced Scorecard: EGI developed strategic-level metrics that are aligned with the EGI Strategic Plan. The key areas of the strategy and the project level objectives were aligned to the metrics using an "EGI Balance Scorecard" which sets out a strategic management and measurement framework that can be used to track the execution of the strategy. The balanced scorecard acts as a measurement system, strategic management system and communication tool. This scorecard will be filled in and revised over the course of the next year.

¹³ <http://go.egi.eu/smc>



Update of EGI Sustainability Plan: Based on the evolution of previous sustainability plan reports, and the work that will take place over the next year and information gathered, a final report will be produced to articulate EGI's updated sustainability plan.

3.3.2 Service Management

If EGI is to continuously evolve as a sustainable service provider, management of those services will need to continuously improve as well. Better service management will offer more predictable service delivery and more efficient use of organisational resources. EGI.eu is working with expert consultants from the FedSM project¹⁴ for improving service management within EGI based on ITIL best practices. Increasing ITSM maturity will continue to be a focus for EGI moving into the future. Specifically though:

Evolve the EGI.eu Service Portfolio: EGI Global Tasks were recently reviewed based on a develop service portfolio¹⁵ and was used as a foundation for evaluating EGI sustainability plan and long-term strategy. Over the next year, this will need to evolve alongside strategic decisions.

Develop the EGI Partnership Service Portfolio (including NGI technical services): The service portfolio will also expand to include the NGI technical services that are user facing so to build a wider service definition covering the EGI partnership.

Improve Service Management: EGI will continue to increase the maturity of its service management processes in the areas of operations, policy and software delivery.

Conduct updated Cost Analysis: An updated cost analysis will also be conducted based on the service portfolio.

3.3.3 Business Development

An essential part of EGI's strategic plan is to ensure that funds for the routine operation of the infrastructure are sustained from stable sources (preferably from within the EGI Community) in order to ensure current and future users of the infrastructure of its sustainability. The funds needed to change and expand the service portfolio will be sought through fixed-term projects, after which the enhanced service portfolio will have to be sustained through community funds and/or business opportunities. The following activities highlight these focus areas:

Form strategic partnerships through MoUs: Strategic partnerships for in-kind effort around areas of mutual benefit (e.g. software, consultancy).

Explore Pay-for-Use and other business models: Through a dedicated pilot group set up to investigate pay-for-use models to adapt to different funding models and/or researchers without access to own resources.

Develop potential professional services: Professional services for delivering specific services from the portfolio to research infrastructures, or expert technical consultancy to an external independent project.

¹⁴ <http://www.fedsm.eu/>

¹⁵ See also MS123 EGI Global Tasks review, <https://documents.egi.eu/document/1566>



3.3.4 Policy Development and Support

Policies are needed to govern the provision of a high-quality distributed-computing production-oriented infrastructure. EGI.eu provides management of policy groups for developing and approving policies relating to operations, software quality, security, user communities and general governance. The development of these policies is coordinated and supported, which may have relevance and impact with other European and International e-Infrastructure providers, therefore, liaison with external groups and organisations is essential as well.

Manage the policy development process: From concept to implementation, the formulation of policies and procedures through EGI policy groups needs to be led, developed and supported. SPT will continue to support the approval process of policies and procedures through the EGI governance.

Liaison with international policy bodies: EGI.eu partners work together with several international policy bodies to establish relationships for European and International policy development. This includes, but not limited to, DG CONNECT through their events (e.g. e-Concertation), the e-Infrastructure Reflection Group (e-IRG), the Open Grid Forum (OGF) e-IRG, EUGridPMA, and IGTF. EC policy also needs to be followed in order to ensure policies are inline with on-going and future developments (e.g. EU2020).

Coordinate policy development and implementation: Coordination of policies (e.g. security) and policy implementation projects such as scientific publications repository, scientific disciplines classification adoption and pay-for-use experiment. This includes work with resource providers for potential policy integration or modification and liaison with groups for specific policies that need to be established or modified for providing services or integrating services.

3.3.5 EGI Compendium

The EGI Compendium is a body of knowledge updated on an annual basis about NGIs/EIROs part of the European Grid Infrastructure. The collected data refers to key areas such as legal form, users, services, resources, technology and sustainability. EGI will collate update information and produce a new edition (2013) in the first half of 2014.

4 VIRTUAL RESEARCH ENVIRONMENTS

Virtual Research Environments (VRE) are defined as the complete and inclusive work environment that is owned, deployed, managed and used by one or more closely related research communities. This definition includes ICT resources that are entirely remote and external to EGI as well as EGI resources that are, or will be, integrated into potential VREs (see Figure 4).

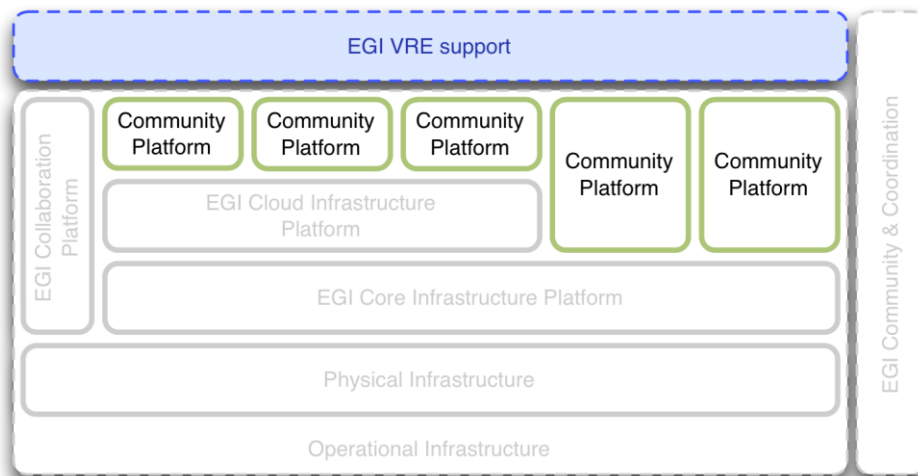


Figure 4: Virtual Research Environments integrate with the EGI Operational Infrastructure through deployed Community Platforms.

This includes the Community Platforms even though they are deployed in the EGI production infrastructure. However, unlike the remainder of the Operational Infrastructure, the Community Platforms are ‘owned’ by the respective Research Communities that are part of the Virtual Research Environment pillar of the EGI ecosystem.

Community Platforms that are deployed on top of EGI’s Cloud Infrastructure Platform are typically deployed and operated by the owning Research Communities, although these operational services may be delegated to EGI. This is accomplished by Research Communities (or their authorised delegates) publishing new VM images and updates to existing images using a VM Image management infrastructure that is similar to a content delivery network, that reaches out to the federated Cloud providers for them to acquire the necessary set of images to support a research community.

This is not true for Community Platforms that integrate directly with the EGI Core Infrastructure Platform: Those Community Platforms *must* be deployed and operated by EGI on behalf of the owning Research Communities. The main integration points are the EGI Software Repository (in form of the UMD) for Resource Centres to use as the main software repository for deployment in the physical infrastructure, and the Software Provisioning Process, which provisions software updates of Community Platforms up to and until their general availability in the EGI Software Repository.

The following subsections summarise the future plans and activities around Virtual Research Environments.

4.1 Community Platforms

As illustrated above, Community Platforms include those IT infrastructure components of a VRE that are deployed remotely on EGI’s production infrastructure. The following describes the activities targeted at Community Platforms.



4.1.1 Unified Middleware Distribution (UMD)

The Unified Middleware Distribution was conceived at the beginning of the EGI-InSPIRE project as a single one-stop-shop repository of software for federated Resource Centres to deploy on their physical infrastructure. The presumption at its inception was that it would be populated with software coming from EGI's main Technology Provider, the EMI and IGE projects, in a truly unified form, i.e. unified interfaces facing the presentation/user layer, and probably a reduction in key components and overall simplification.

As it turned out, the level of integration and uniformity was less than expected, which resulted in the UMD repository basically containing a number of parallel existing Community Platforms, provided by the EMI and the IGE projects. While the key distribution milestones for the IGE project, the releases of IGE-1, IGE-2 and IGE-3 represent snapshots of an otherwise continuous evolution of the IGE Community Platform based on the Globus Toolkit and services, the EMI project introduced a number of backwards-incompatible changes in EMI-2 and EMI-3, thus effectively producing three very similar yet distinctively separate Community Platforms: EMI-1, EMI-2 and EMI-3.

The following activities are planned for the further maintenance of the UMD:

Continue providing UMD releases: For as long as the EMI consortium is publishing coordinated updates to the EMI-2 and EMI-3 Community Platforms, EGI will continue provisioning these into the corresponding UMD-2 and UMD-3 distributions. The same will hold true for updates that may be provided by the IGE consortium. PSNC will continue to provide updates for the QCG platform to be made available in future updates to UMD-3.

Separation of Community Platforms: Currently, all existing Community Platforms are collated in the UMD repository. With the end of the IGE and EMI projects, significant coordination efforts will no longer exist. Lack of cross-community platform coordination will result in a significantly higher probability of severe disruption of software updates due to incompatibilities in libraries and shared code. Therefore, Community Platforms must be further separated hence protected from each other to manage the risk of EGI service disruptions. While the overall software process may not change significantly, the output target for each provisioned Community Platform will not be a shared UMD repository, but individual repositories that in their entirety constitute the "new UMD".

4.1.2 Emerging Community Platforms

The software bundled into the UMD was primarily used by those Research Communities known as EGI's "Heavy User Communities": researchers and research groups coming from the High Energy Physics, Life Science, Astronomy & Astrophysics, and the Earth Sciences were and still are the primary users of the EMI Community Platforms.

However, with the end of the EMI and IGE projects, it is expected that coordinated EMI releases will disassemble into individual Community Platforms around ARC, UNICORE and dCache as sustainable middleware stacks, while the IGE platform around the Globus Toolkit is likely to remain the same. At the moment the future of the components of the former gLite middleware is unclear.

Thus, EGI expects the following Community Platforms to emerge, which will require to be deployed directly on the physical infrastructure:

UNICORE HPC: Traditionally serving the HPC community, the UNICORE Community Platform remains stable yet receiving steady updates.



ARC: The Nordic academic community has grouped itself around the ARC Grid middleware and will continue to support it beyond the EMI project.

dCache: The dCache platform is a scalable storage platform that in and by itself serves a number of use cases up to and including data oriented infrastructures.

EGCF Globus: The European Globus Community Forum (EGCF) has constituted itself as the successor of the IGE project. Members of the EGCF will take care of a sustainable maintenance activity around providing and maintaining a Globus-based Community Platform to interested Research Communities.

QosCosGrid: Currently being integrated in to UMD releases, EGI expects that the QCG platform will continue to be deployed in the EGI production infrastructure.

4.2 Support Services for VREs

The PY4 plan of the EGI.eu User Community Support Team follows up the main areas of work from PY3, but with more emphasis on two activities:

- Reaching out to the long tail of researchers and get them on board of the federated cloud or existing grid based infrastructure.
- Federate resources from large research communities through the core operational tools.

Activities supporting Virtual Research Environments are grouped by priority according to their applicability and contribution towards EGI's goals for H2020:

High-priority activities:

Targeted projects: Establish and run focussed projects for the benefit of scientific groups in the ERA. This includes existing and new Virtual Teams and mini projects (see section 3.1.5), pilot activities, case studies, etc.

Cloud-based VREs: Increase the uptake of the EGI Cloud Infrastructure Platform by community-specific VREs and platforms (sharing tools, technical documentation, best practices)

Community engagement: Technical engagement with new communities of the ERA by attending and contributing to their and to joint events (attend by NGIs and/or EGI.eu)

NGI International Liaisons & EGI Champions: Maintain and develop the network of NILs and Champions. Coordinate their technical outreach to ESFRI project member institutes.

Low-priority activities:

Bare-metal based VREs: Increase the number of Virtual Research Environments that directly integrate with the EGI Core Infrastructure Platform.

Collaborative VREs: Expand the EGI Collaboration Platform with new services that simplify the use of the EGI production infrastructure by communities and their VREs. Generic (enough) higher-level



services to access and consume the EGI production infrastructure are candidates for inclusion in the EGI Collaboration Platform. Candidates that are currently evaluated are DIRAC and the SHIWA Portal.

Untargeted projects: Establish and run projects (e.g. VT projects or pilots with the NGIs) that have no identified research group beneficiaries in the ERA.

Unfunded best-effort activities:

Untargeted evaluations: Evaluate technology that has no identified direct beneficiaries in the ERA.

Untargeted dissemination: Contributing to dissemination events that have no specific target community in the ERA.



5 CONCLUSIONS

Supported by comprehensive information provided in recent technical documentation, this roadmap summarises the plans around activities and improvements for the last project year of the EGI-InSPIRE project. Structured around EGI's three strategic pillars for Horizon 2020, this document further organises the available information according to the technical architecture of the EGI operational infrastructure.

The profile of the Core Infrastructure Platform is sharpened towards a toolkit with surrounding services that can be offered and taken up to federate any distributed infrastructure into a consistent and uniform delivery of services. The Cloud Infrastructure Platform is maturing from a prototype into a potential service offering that may be put into production in the mid-term future, while at the same time serving as a reference platform for integrating with the Core Infrastructure Platform. The Collaboration Platform settles into a set of services and solutions that allow research communities to integrate EGI resources seamlessly into their own extended virtual research environments.

Activities related to communities and coordination complement EGI's activities in that they provide and facilitate human and social coherence and sharing of knowledge across the whole EGI ecosystem. Covering human networks, communication & marketing, and policy & strategy these activities provide necessary cohesion for internal communication (human networks), external communication (communication and marketing) and general direction of evolving EGI (strategy and policy).

Support for Virtual Research Environments include both technical support and consultancy to research communities. These cover help with defining the purpose and contents of Community Platforms, as well as engaging with research communities with the clear goal of those communities using the EGI infrastructure in the medium term.

6 REFERENCES

R1	D2.21 Annual Report on EGI and its External Relations Activity https://documents.egi.eu/document/1578
R2	D2.22 Marketing and Communication Plan https://documents.egi.eu/document/1762
R3	D4.8 Annual Report on the EGI Production Infrastructure https://documents.egi.eu/document/1664
R4	MS425 Operational Level Agreements (OLAs) within the EGI production infrastructure https://documents.egi.eu/document/1712
R5	D5.9 Annual Report on the status of Software Provisioning activity https://documents.egi.eu/document/1657
R6	MS514 EGI Platforms Roadmap https://documents.egi.eu/document/1624
R7	MS520 EGI Federated Cloud Blueprint V1 https://documents.egi.eu/document/1773
R8	D6.9 Annual Report on the HUC Tools and Services https://documents.egi.eu/document/745
R9	D7.3 Annual Report on Operational Tool Maintenance and Development Activity https://documents.egi.eu/document/1562
R10	MS710 Roadmap for the Maintenance and Development of the deployed Operational Tools https://documents.egi.eu/document/1501