EGI-InSPIRE

EGI TECHNICAL ROADMAP

EU DELIVERABLE: D2.35

Document identifier:	EGI-InSPIRE-D2.35-final.docx
Date:	**17/06/2014**
Activity:	**NA2**
Lead Partner:	**EGI.eu**
Document Status:	**FINAL**
Dissemination Level:	**PUBLIC**
Document Link:	https://documents.egi.eu/document/2207

Abstract

This deliverable constitutes the third edition of the EGI Technical Roadmap. It provides an executive overview about the project's planned activities for project year 5. 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	EGI.eu / NA2	18 June 2014
Reviewed by	**AMB & PMB**		17 June 2914
Approved by	**AMB & PMB**		18 June 2014

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	3 Jun 2014	Initial, internal version	Michel Drescher, EGI.eu
2	4 Jun 2014	Nearly complete section 1, 2	Michel Drescher, EGI.eu
3	5 Jun 2014	Completed section 4	Michel Drescher, EGI.eu
4	6 Jun 2014	Completed section 2, 3.1, 3.2 & 5	Michel Drescher, EGI.eu
5	9 Jun 2014	Final draft	Michel Drescher, EGI.eu
6	12 Jun 2014	AMB review revision	Michel Drescher, EGI.eu
7	17 June 2014	Added missing metadata	Michel Drescher, EGI.eu

IV. APPLICATION AREA

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

V. DOCUMENT AMENDMENT PROCEDURE

Amendments, comments and suggestions should be sent to the authors. The procedures documented in the EGI-InSPIRE "Document Management Procedure" will be followed:
https://wiki.egi.eu/wiki/Procedures

VI. TERMINOLOGY

A complete project glossary is provided at the following page: http://www.egi.eu/about/glossary/.

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 third edition of the EGI Technical Roadmap. It provides an executive overview about the project's planned activities for project year 5. Where applicable, interactions and collaborations with external projects are briefly mentioned.

The first edition began structuring its content around the three pillars of the EGI strategy, namely: (1) Operational Infrastructure, (2) Community & Coordination, and (3) Virtual Research Environments. The second edition [R 1] continued 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.

This third edition continues with the established document structure and purpose, while introducing the changes that are either planned for PY5 or emerged unexpectedly in the past.

The EGI production infrastructure continues to be maintained as set of well-defined platforms. However, EGI attention to Community Roadmaps is focused around the actual release process and details of technical integration with EGI platforms. The Technical Roadmap will refer to Community Platforms in terms of inclusion or exclusion as part of the roadmap of the UMD, but no further architectural decomposition or planning activities are conducted.

The roadmap includes milestones from March 2014 onwards.

TABLE OF CONTENTS

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.

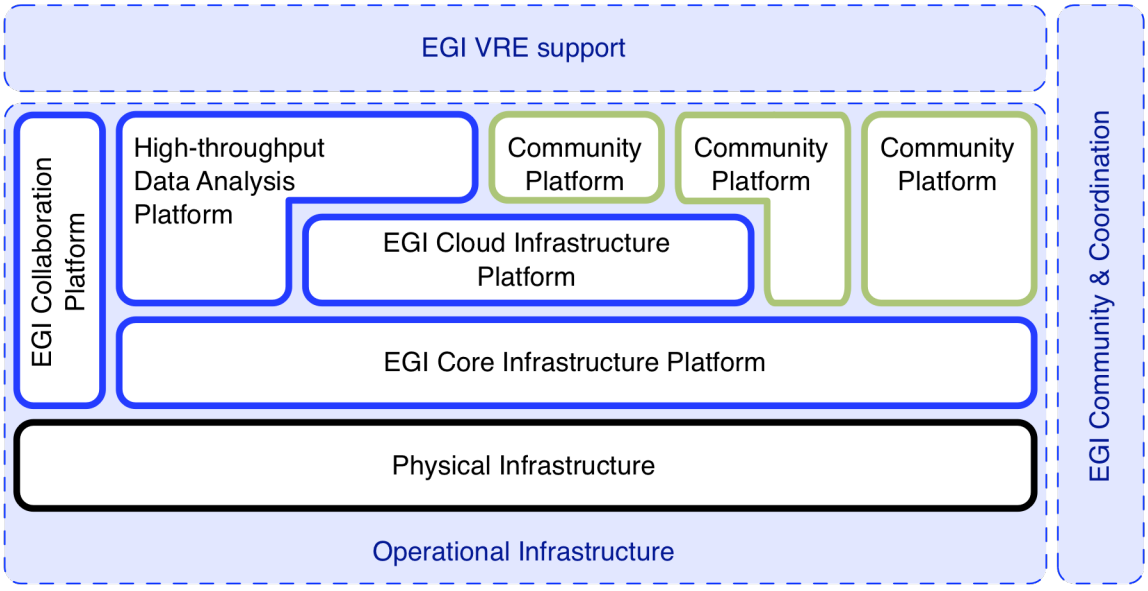

Figure 1: The three pillars of the EGI strategy

Receiving annual updates, the EGI Technical Roadmap is based on its previous editions 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.

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 four 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.

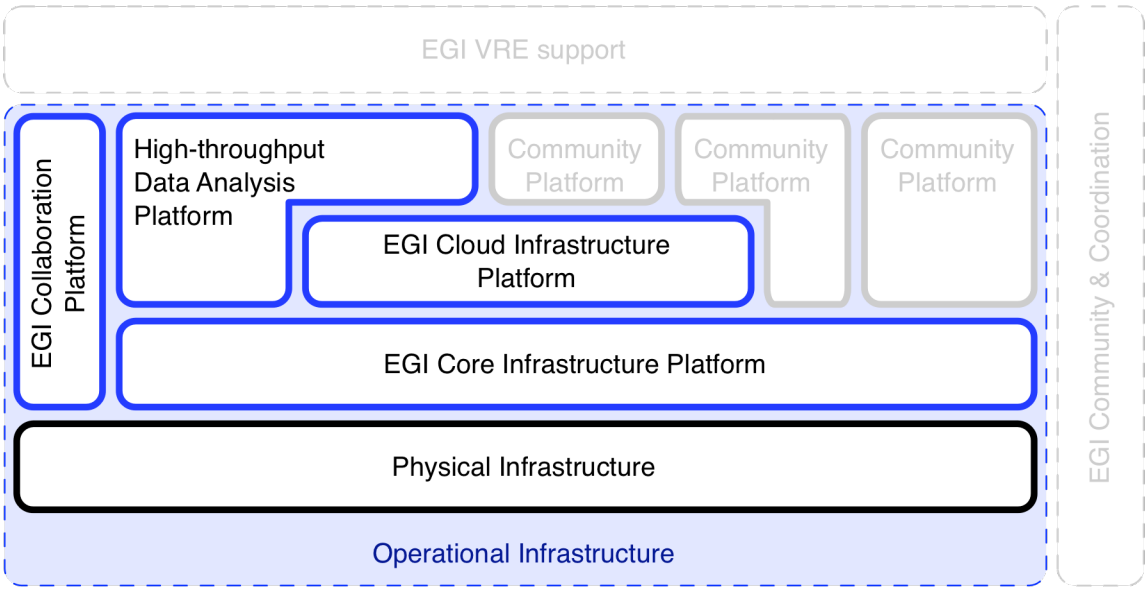

Figure 2: The first pillar of the EGI Strategy is organised in four distinct platforms

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, High-throughput Data Analysis 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.

Currently, there are no significant changes to the Federated AAI foreseen in PY5.

2.1.1.2 Information Discovery

The Information Discovery subsystem comprises of local BDII instances deployed at every federated Resource Provider, and a number of top BDII instances that aggregate information gathered from local BDII instances (as per configuration) for a partial-global view of the production infrastructure for programmatic enquiry.

Currently, there are no significant changes to the Information Discovery Subsystem foreseen in PY5.

2.1.1.3 Monitoring

The Service Availability Monitoring (SAM) subsystem provides EGI with a comprehensive repository of availability and reliability information for each federated Resource Provider.

Over the coming months, the current SAM framework will undergo a major overhaul, effectively replacing all existing components. This includes the integration of the ARGO availability and reliability computation component that ahs been developed in the mini project TSA4.10. Details for each activity can be found in MS711 [R 2].

Activity / Task	Planned completion time
Migrate central SAM operations to new consortium (CRNS, SRCE, GRNET)	May 2014
Development of a new web user interface replacing MyEGI	December 2014
Extension of the WebAPI delivered by mini project	December 2014
The extension of the Sync Components	December 2014
Creation of a leaner monitoring instance	December 2014
Removal of the Oracle database dependency	December 2014

Integrate ARGO (from TSA4.10) as beta component in SAM	December 2014
Switch to ARGO & end of life of ACE	April 2015

2.1.1.4 Accounting Repository

The Accounting Repository is a central service aggregating individual accounting data sent from local Resource Providers using the SSM (secure stomp messenger) protocol version 2 via the EGI Messaging network.

For PY5 the following activities and tasks are scheduled. Details for each activity can be found in MS711 [R 2].

Activity / Task	Planned completion time
Cloud Accounting to production	April 2014
Storage Accounting Summaries to Accounting Portal	June 2014
Send MPI data to portal	June 2014
Application accounting prototype	June 2014
Application Accounting usage record finalised	September 2014
Publishing summaries from Accounting Repository to other sites (OSG/DGAS)	September 2014
Send all accounting data (including MPI data) to portal in production using the new schema	September 2014
Migrate sites from SSM1.2 to SSM2	December 2014
Storage Accounting in production	December 2014
GPGPU accounting prototype	December 2014
Improvement of the cloud accounting to cover storage accounting for transient cloud storage and data usage accounting by the virtual machines	December 2014
Adoption the OGF Usage Record v2	December 2014
Support to implement Pay-for-Use proof of concept	December 2014

2.1.1.5 Accounting Portal

Although by capability a component of the Accounting subsystem, the Accounting Portal is developed by a different EGI federation member hence accounted for and managed in an independent roadmap section. For PY5 the following activities are scheduled. Details for each activity can be found in MS711 [R 2].

Activity / Task	Planned completion time
SSM implementation for CPU Accounting	June 2014
XML endpoints generalization and improvement	June 2014
Provisioning of Application accounting (DB implementation)	October 2014
Provisioning of Application accounting (View implementation)	October 2014
Scientific Disciplines VT Interface Support	October 2014

Scientific Disciplines VT final Implementation	October 2014
Adoption the OGF Usage Record v2	October 2014
Support to implement Pay-for-Use proof of concept	November2014
Regional Portal Implantation in other NGIs	December2014
Improvements Storage View	December 2014
Improvements Cloud View	December 2014
Improvements MPI View	December 2014
General Improvements	December 2014
Portal extensions to include GPGPU usage information	December 2014

2.1.1.6 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.

There are currently no significant changes foreseen for the messaging infrastructure.

2.1.1.7 Central Service Registry

During PY5 new features for GOCDB will be developed outside of EGI as an open source project. Future developments are therefore largely undetermined and are likely to evolve. Nevertheless, continued involvement within the EGI Global Task will help ensure future developments are strategic and interoperable. Support will also be provided in PY5 to provide continued EGI operational support, service hosting and bug fixing. The major developments for GOCDB are likely to include multiple service endpoints and further GLUE2 support.

For PY5 the following activities and tasks are scheduled. Details for each activity can be found in MS711 [R 2].

Activity / Task	Planned completion time
GLUE2 XML rendering of GOCDB data	Aug/Sept 2014
Writable PI method to submit downtimes	~June 2014
Multiple Service Endpoints	~May 2014
Extend data model and add more GLUE2 attributes (e.g. GLUE2 cloud extensions)	~July 2014
Web portal interface enhancements	To be defined

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

rota[1]. We plan though to expand security-monitoring services to reflect the current needs especially to support more extensive deployment of the patch monitoring and explore possibilities to collect security-related characteristics from IaaS clouds.The activities for PY5 are scheduled as follows.

Activity / Task	Planned completion time
Packaging Pakiti client for Linux distribution, integration with EPEL (if possible)	September 2014
Report on possibilities for security monitoring in IaaS clouds	November 2014

2.1.2.2 Operations Portal

The Operations Portal is one of EGI's key tools to enable and facilitate federated operation of a federated distributed production infrastructure. It provides key dashboards for regional operator on demand (ROD) roles and other resource and VO related information management.

For PY5 the following activities and tasks are scheduled. Details for each activity can be found in MS711 [R 2].

Activity / Task	Planned completion time
Provide packages that are easy to install via platform package management	September 2014
New disciplines classification	June 2014
Improving of the Operations Portal v3 according to the users feedback	July 2014

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 has matured over time, and is very stable and productive in its outcome.

The same process is applied to every Community Platform that is included in the UMD, with the same procedures, criteria and performance indicators. The process was generalised over time so that differences across Community Platforms are aggregated in platform-specific documents (e.g. applicable Quality Criteria may differ from platform to platform).

All necessary sub-processes and activities (i.e. Quality Assurances (Quality Control, Staged Rollout, and finally UMD publication in the EGI Software Repository) are integrated and coordinated through regular, typically weekly UMD Release Team conference calls.

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 both public and private IaaS Clouds. It wholly embraces the Cloud paradigm and extends it with a federation mechanism that is partially

[1] https://wiki.egi.eu/csirt/index.php/Security_Officer_on_Duty_tasks

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 becoming fully integrated into the EGI Production Infrastructure following its launch in May 2014 with a core set of user-facing service and capabilities that will be extended over time.

2.2.1 Federated Cloud management services

2.2.1.1 VM Management

VM Instance Management is provided by Resource Providers exposing an OGF OCCI compliant interface for consumption by the user. As of date this is currently version 1.1. The actual implementation of the defined functions is left to the Resource Provider to source in; this is typically provided by either the core Cloud Management Framework's (CMF) components, or by fronting services (e.g. the rOCCI-server).

The following activities are foreseen for PY5:

Activity / Task	Planned completion time
Complete integration with Information Discovery System	July 2014
Extend network management features according to OCCI spec	September 2014
Push agreed contextualisation profile for OCCI into OCCI WG, including possible extensions for a possible OCCI v2 specification	December 2014

2.2.1.2 Data Management

Standards-based Cloud Storage and Data management in EGI has a very similar architecture as the VM Instance Management service: A proxy-server is fronting actual Cloud Management Frameworks that do not natively support CDMI v1.0 (the chosen standard in EGI) as their access interface. The prototype implementation has been further refined in PY4 though its deployment and widespread use in production had to be delayed due to external factors beyond the project's control.

Activity / Task	Planned completion time
Deploy CDMI-based proxy service for OpenStack in testbed	August 2014
Support and maintenance of any detected bugs and omissions	December 2014

2.2.1.3 Image management and distribution

VM Images, once registered in the collaborative Cloud Marketplace (see Application Database in Collaboration Platform), need to be distributed to the participating Resource Providers. This includes managing updates applied to images and image lists. The tool to support the distribution of the images from their repositories or other locations is the HEPIX VMCaster[2] product. This has been integrated with the Application Database and also packaged for easy deployment within the Cloud management frameworks

Currently there are no significant improvements foreseen for this component though participants of the federated cloud group will act on support requests.

[2] https://github.com/hepix-virtualisation/vmcaster

2.2.2 EGI Platform integration activities

As described in MS518 [R 3] the Cloud Infrastructure Platform makes extensive use of the Core Infrastructure Platform and the Collaboration platform. To improve clarity for the reader, this subsection of the roadmap will collate all technical activities across the EGI Platforms that are related to the integration of the Cloud Infrastructure Platform as a fist-class member of the EGI production infrastructure.

2.2.2.1 Virtual Organisation management & AAI

The Federated Cloud adopts the Core Infrastructure's approach to federated AAI. In that respect, any changes in the Core Infrastructure automatically will affect the Cloud Infrastructure Platform.

In itself, AAI integration is accomplished; however, with the recent development and further deployment of EGI's e-Grant tool, the following activity is foreseen in PY5:

Activity / Task	Planned completion time
Feasibility investigation to integrate Perun VO management with automatic resource allocation through e-Grant	August 2014

2.2.2.2 Monitoring

The federated Cloud resources will be monitored through EGI's existing SAM subsystem.

The following activities for PY5 are foreseen:

Activity / Task	Planned completion time
Finalise deployment of OCCI interface monitoring probe	June 2014
Finalise CDMI interface monitoring probe, including production deployment	July 2014
Broker services monitoring probes deployed in production	December 2014
Benchmarking of connected Cloud services will be investigated with available tools from the open market	December 2014

2.2.2.3 Accounting

The EGI Cloud Infrastructure Platform makes use of the Accounting subsystem of the Core Infrastructure Platform. The following activities and improvements are foreseen for the accounting of Cloud resource usage in PY5:

Activity / Task	Planned completion time
Deploy SSL/TLS support for Cloud accounting data	June 2014
Audit and verify reported network, memory & disk usage figures for Cloud resources	August 2014
Integration with pay-for-use activity in EGI • VM size related charging mode • Charging for CPU hours vs. wall clock time The work generally largely depends on local Resource Provider polity,	Ongoing, in very early stage

2.2.2.4 Information Discovery

Information Discovery in the EGI Cloud Infrastructure Platform provides a mix of static and semi-static technical information about a resource provider's service offering, either as complimentary information to the baseline federated information (for example, which Cloud Management Framework and version of same 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).

It re-uses the Core Infrastructure Platform Information Discovery subsystem based on BDII; however many Cloud related information discovery elements require close collaboration and extension with the Core Infrastructure product management.

The following activities for PY5 are foreseen:

Activity / Task	Planned completion time
Migrating all Resource providers to latest LDIF schema	June 2014
Investigating requirements for dynamic information in local BDII • Grouping images and image lists per VO • Listing public and private images and image lists • Current resource allocation load vs. free resources • Available SLAs	December 2014

2.2.2.5 Central Service Catalogue

All services and service endpoints of the EGI Cloud Infrastructure Platform are registered in EGI's central service catalogue. This service is sufficiently generic to not require substantial changes for PY5.

2.2.2.6 Operational Security

Operational Security support services deal with the containment of discovered infrastructure compromisations (CSIRT), handling software vulnerabilities (RAT), preventative security measures and assured, secure certification of Resource Providers for production integration.

For PY5 the following activities are foreseen:

Activity / Task	Planned completion time
Update RP certification PROC18 to include security questionnaire provisions	July 2014
Update RP certification PROC18 to include a simple security challenge assessing RP's operational compliance with EGI basic security policies	December 2014

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. For PY5 the following activities and tasks are scheduled. Details for each activity can be found in MS711 [R 2].

Activity / Task	Planned completion time
Implementation of alarm processes for EGI tools	March 2014
Additional authentication through shibboleth	March 2014
New interfaces to PRACE and XSEDE	Depend on PRACE and XSEDE
CMS specific adaptations	March 2014
Merge GGUS and xGUS webfrontends to a common platform	March 2014
Implement a bulk submit feature to enable the notification of many sites at the same time	March 2014

2.3.1.2 VM Marketplace & Appliance Repository

The VM Marketplace and Appliance repository are discontinued. As a replacement, a Cloud marketplace has been developed and integrated into the Application Database.

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 Software Repository has matured into its originally desired state. No more activities are planned beyond maintenance of the current functionality.

2.3.1.4 Application Database

The EGI Applications Database (AppDB) is a centralised service that stores and provides information to EGI members, and to the general public about:

- Tailor-made scientific applications that are integrated with the EGI production infrastructure, or with some EGI partner infrastructure (for example with a desktop grid).
- Software tools, components and frameworks that application developers can use to integrate new scientific models and applications with the EGI production infrastructure, or with some EGI partner infrastructure.
- Publications about the aforementioned scientific application and software items.
- Programmers and scientists, who develop, drive the development and/or provide user support about the above software.

During the next year the further development of AppDB is expected to be driven by the needs of federated cloud users. One one hand this will require new features/services in AppDB that improves the workflow of exchanging Virtual Machine Appliances among users and user communities, and on the other hand enables and makes even more convenient to deploy and instantiate Virtual Appliances on federated cloud sites that participate in EGI.

Activity / Task	Planned completion time
Full availability of the Cloud Marketplace for public VM image lists	June 2014
Full support for private image lists for VOs and or individuals	June 2014
Automatic image subscription for Resource Providers enabling a given VO	December 2014

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 Training Marketplace is an online registry to advertise and to view (browse and search) training events, online training materials, training resources and university courses that relate to EGI. The service supports cooperation between trainers and trainees in different localities and projects by connecting the groups through the stored items that are advertised in the Training Marketplace. The Training Marketplace is typically used together with the EGI Document Database to index and highlight materials files from the Document Database that relate to training. The Training Marketplace provides rating and commenting facilities the registered items, and web gadgets for integration with third party websites, such as NGI sites and research community sites.

The service has reached its desired maturity and feature set; starting from April 2014 this activity will receive no more EGI-InSPIRE funds. However the current service provider, STFC, will maintain this service as-is for EGI until the end of the EGI-InSPIRE project in December 2014.

2.3.2.2 Client Relationship Management

EGI's Client Relationship Management System is a client database with a web interface providing capabilities for the NGIs to record contact leads to new communities and e-infrastructure requirements captured through these leads about the new communities. The CRM system also provides statistical overviews and reports about the community engagement activities and therefore it helps EGI gain a global understanding of the needs of new communities, as well as the effectiveness and performance of the NGIs' outreach and engagement activities.

Operating the CRM system will discontinue beyond April 2014. The database content will be preserved for possible future use, and the software is archived in a virtual machine image to be re-activated for a possible future use on the EGI Federated Cloud infrastructure.

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.

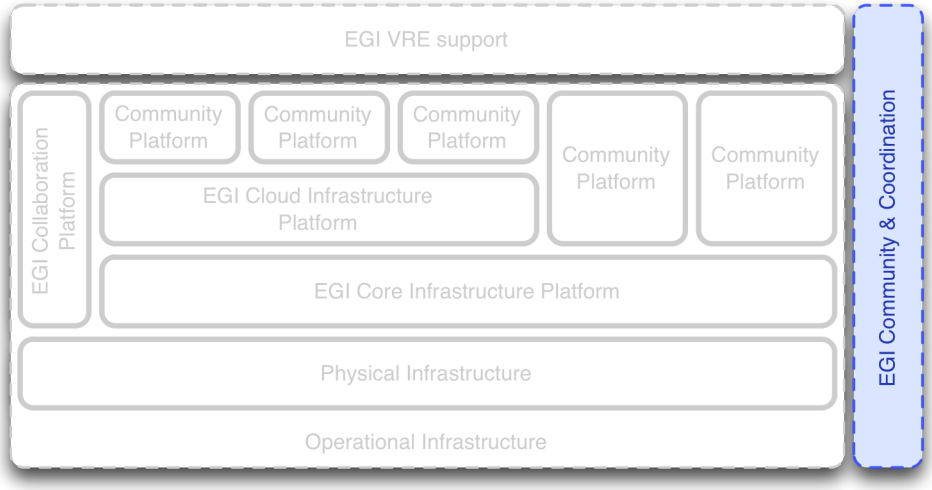

Figure 3: EGI Community and Coordination services form the second pillar of the EGI strategy.

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

3.1.1 Distributed Competence Centres

During PY4 another human network, the Distributed Competence Centre[3] has been established to support the technical engagement with new communities. The DCC includes skilled user-support personnel and technical assets that can be accessed by the EGI community to support the uptake of EGI services by new user communities and Research Infrastructures. The DCC works as a distributed team of experts run under the EGI.eu coordination. Experts from the DCC are appointed for to help EGI capture, refine and document requirements of specific new communities or in other word to support the scoping of new Virtual Team projects. DCC members are also able to join Virtual Team projects and help the implementation with technical knowledge about certain tools/software/system that is in the topic of the Virtual Team.

During PY5, EGI will conduct an open call for DCC that will be further evolved through follow-on projects to EGI-InSPIRE.

3.1.2 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.

Coordination and support of the NGI International Liaisons network is planned to continue without major change in PY5 and beyond, except for merging the UCB, NIL and EGI Champions recurring teleconferences into one stream of then slightly larger conference calls, recognising the benefits of direct information exchange between these groups, and the strong links between the different roles in some of the countries.

[3] http://go.egi.eu/dcc

3.1.3 EGI Champions

The EGI Champions programme establishes a cadre of 'ambassador' scientists and researches, acting as enthusiastic and proactive promoters of EGI towards its respective scientific communities and beyond.

The programme ahs been established and operational since a while now, and no major changes are foreseen, except that the to date separate EGI Champions conference calls will be merged into a joint UCB, NILs and EGI Champions conference call (see above).

3.1.4 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.

During PY5 there are no changes foreseen to the current Operations network.

3.1.5 Virtual Teams

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.

While there will be no more funded mini projects in EGI-InSPIRE PY5,the VT scheme is now a stable and sustainable framework for lightweight community-driven improvements of the EGI infrastructure:

- The EGI User Community Support Team will continue to improve the process and framework through providing template report documents, procedures, and assistance in collaboration available tools (e.g. WebEx, Indico)
- EGI User Community Support Team will introduce more agile project management methodologies to the framework, to respond to the challenges of managing projects with unfunded members and targets that become explored during the project itself.

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.

The roadmap for PY4 has been implemented, and the activities in this area are going into maintenance; no significant changes are foreseen for PY5.

3.3 Strategy, Policy and Business Development

Over the course of the second part of 2014, the activities will be devoted to consolidating the updates on the EGI strategy into a new document and align the balanced scorecard for measuring the progress. The planned activity on the proof of concept for a pay for use provision of EGI services will be also

completed with the writing of a final report documenting the technical, organisational and legal aspects. The activity will also be focused on finalising a business engagement program for SMEs & industries.

Activity / Task	Planned completion time
Update the EGI Strategy document	September 2014
Complete the EGI partnership portfolio	October 2014
Develop business engagement program with SME & Industry	October 2014
Complete pay for use proof of concept	November 2014
Align the EGI Balanced Scorecard to strategy update	December 2014
Support evolution of IT service management best practices	December 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).

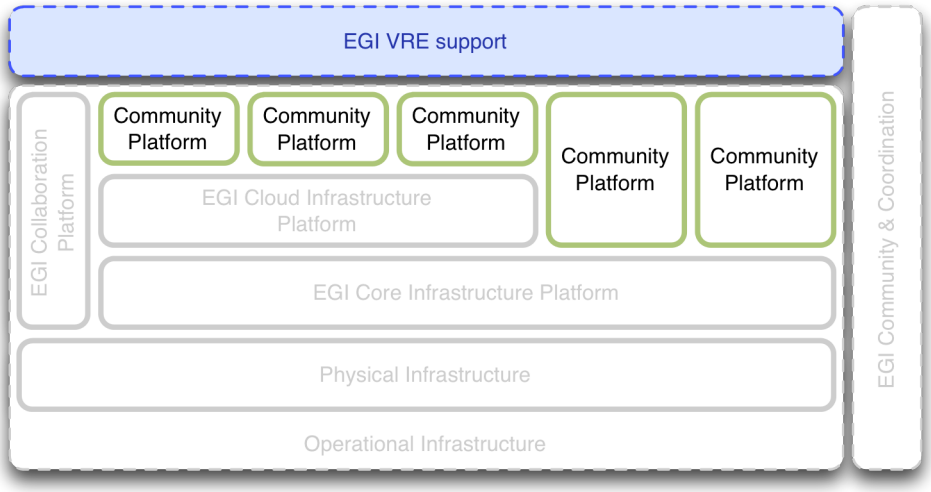

Figure 4: Virtual Research Environments integrate with the EGI Operational Infrastructure through deployed Community Platforms.

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 (c.f. to the EGI Platform Roadmap [R 3]). 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 (see also EGI Platform Roadmap [R 3])

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. By legacy the UMD turned out as a collection of relatively independent yet interoperable Grid Computing platforms. With the IGE and EMI projects ending, the sustainability of the UMD and associated platforms had to be reconsidered.

The result is a UMD that hosts a number of independent Grid Computing community platforms, alongside an emerging EGI High-throughput data analysis platform (see EGI Platform Roadmap [R 3]) built using a reduced set of components coming from the IGE and EMI projects.

In any way, the UMD and the corresponding Software Provisioning process have matured into stable maintenance-only components of the EGI services.

4.1.2 Community Platforms

As expected, the integrated software releases provided by IGE and EMI disassembled into relatively independent Community Platforms that already existed before the EMI project.

This section is providing a brief overview of the Community Platforms that are currently hosted in the UMD. Next to these, a variety of community-managed applications and computing platforms are available in the Community Repository, managed online[4] through the Application Database service without curation or coordination conducted by EGI.

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.

[4] https://appdb/egi.eu

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

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.

Most activities in the EGI-InSPIRE landscape are matured, and operate as designed. Naturally, the scope and size of new developments is greatly reduced compared to the activities that were taking place over the last year.

The one exception is the EGI Cloud Infrastructure Platform. With it becoming integrated with the EGI production infrastructure starting in May 2014, a newly launched first-class resource infrastructure naturally receives most of the attention and effort for improvements and integration effort.

6 REFERENCES

R 1	D2.33 EGI Technical Roadmap 2nd edition, https://documents.egi.eu/document/1706
R 2	MS711 Roadmap for the maintenance and development of the deployed operational tools, https://documents.egi.eu/document/2069
R 3	MS518 EGI Platform Roadmap, https://documents.egi.eu/document/2232