

E-INFRASTRUCTURE INTEGRATION PROPOSAL

Working document to support the development of a joint vision

Authors: T. Ferrari/EGI.eu

Date: 02 Dec 2012

Document status: **DRAFT**

Document link: <https://documents.egi.eu/document/2051>

1 Executive Summary

This paper presents five different collaboration opportunities in the areas of service provisioning, operations, user services and sustainability:

1. Integration of the core infrastructure services
2. Full integration of Common Data and Execution Services
3. Integration of generic operational services and activities
4. Integration of user-facing services
5. Business development and sustainability

Collaboration opportunities will be discussed by the e-Infrastructure representatives to establish a common integration roadmap to better serve user communities by providing one front desk and ensuring expertise can be easily made available, foster integration and reuse of common infrastructure services, and sharing of common operational services.

DISCLAIMER. This is a DRAFT document to support discussion during the EGI Horizon 2020 workshop in Amsterdam and will evolve according to the outcome of it.



2 Introduction

According to the e-IRG Roadmap 2012¹ in order to meet the challenges of implementing the 2020 Strategy, “Europe needs an e-Infrastructure Commons for knowledge, science and innovation” which “should be open and accessible, and continuously adapting to the changing requirements of research and to new technological opportunities”.

As stated in the e-IRG White Paper 2013² “no single e-Infrastructure provider currently provides a full portfolio of e-Infrastructure services needed by the end users. However, users want a single and easy to use interface to all e-Infrastructure services they need. They need services that are coherent, managed and above all integrated so that they can get on with the business of science. They also need a constant innovation of these services, way ahead of what commercial providers can offer.”

In addition, e-IRG believes this challenge can be met by maintaining a separation between the three different functions:

1. **Community building, high-level strategy and coordination in Europe:** for each type of e-Infrastructure service, a single coordinating organisation with a central role for user communities. These bodies, in turn, will need a forum for coordination between them across the different e-Infrastructure types.
2. **Service provision:** flexible, open, and competitive approach to national, European, and global service provision; with advanced collaboration among the interested public and commercial service providers.
3. **Innovation:** Implementation of major innovation projects through the best consortia including e-Infrastructure suppliers, industry, users and academia with a dedicated management structure comprising the partners per project.

The e-IRG White Paper 2013 finally recommends that: “International organizations of e-Infrastructures should join forces and share their common challenges towards serving the European user communities, thereby avoiding duplication of efforts (as far as possible) in such areas as:

- Outreach to and involvement of user communities;
- Services registry, discovery and provisioning;
- Financial, legal, business development and procurement issues

This paper identifies the integration areas that EGI see as potential

3 User Communities

In November 2012 EGI, EUDAT and PRACE started a collaboration to identify the interoperability and integration needs of user communities requiring services from different e-Infrastructures. Test environments were setup using existing technologies like GridFTP, UNICORE, Globus Online, iRODS, EMI data management services, etc³. The following is a non-exhaustive list of user communities and Research Infrastructures that require the integration of EGI, EUDAT and PRACE services, and were involved in the pilot activities:

¹ <http://www.e-irg.eu/news/news/505/e-irgs-white-paper-2013-published-europe-needs-a-e-infrastructure-commons.html>

² http://www.e-irg.eu/images/stories/dissemination/white-paper_2013.pdf

³ https://wiki.e-irg.eu/wiki/EGI_EUDAT_PRACE_collaboration



- [Distributed Research Infrastructure for Hydro-Meteorology Study](#) - DRIHMS (hydro-meteorology)
- [European Plate Observing System](#) - EPOS (European plate observation)
- [Multiscale applications for European e-Infrastructures](#) - MAPPER (multi-scale simulation for fusion, hydrology, physiology, nano-material science, computational biology)
- Molecular and materials science
- [Virtual Earthquake and seismology Research Community in Europe e-science environment](#) - VERCE (seismology)
- [Virtual Physiological Human](#) - VPH (biomedical modelling and simulation of the human body)
- The EGI Collaboration and EUDAT are also successfully collaborating to support EISCAT-3D (in the framework of the ENVRI cluster project, LifeWatch and EPOS).

These pilot activities were focused on the technical integration of various service components of EGI and EUDAT for data management, data staging and registration of persistent identifiers.

The following sections provide an overview of the status of technical integration work and identify other areas that require further integration effort.

The following diagram illustrates the EGI platform architecture.

- The EGI Core Platform includes Authentication, Authorization and Identification. Its main scope is to enable flexible and efficient provisioning of IT resources irrespective of the customer's actual use of those resources. The Core Infrastructure Platform will form the foundation layer of all other platforms that are, or will be, built on top of it.
- The collaboration platform is enabling the collaboration between communities that are using technology deployed on top of the core Infrastructure Platform. It builds on top of the existing EGI production infrastructure, so that EGI's research communities are able to transition. The EGI Collaboration Platform comprises services and technology that are (expected to be) used across many if not all EGI research communities. Services for data management, data transfer and replication, and metadata catalogues are examples.
- EGI Community Platforms are best described as meeting the needs of the respective community.




Figure 1. EGI Platform architecture

4 Objective 1: Integration of the core infrastructure services

- A common infrastructure of authentication authorization and identify provisioning is needed to support users accessing services from different services.
- An integrated infrastructure for Permanent Identifiers need to be defined so that data PIDs can be resolved and provided.

None of the two activities was successfully tested in the pilot program conducted so far. Authorization using X.509 certificates required manual intervention on EUDAT sites.

EGI and EUDAT should collaborate for the collection of requirements from service providers for the wider adoption of a federated identify provisioning infrastructure.

5 Objective 2: Full Integration of Common Data and Execution Services (EGI Collaborative Platforms)

Integration of data and execution services for data processing was tested by EGI, EUDAT and PRACE to ensure that data can be seamlessly staged, processed, transferred, referenced across the various infrastructures. The objective this integration activity is to ensure the full interoperability.

Achieved

- Data Access and Transfer (SRM and GridFTP): EGI (SRM – DPM, dCache, StoRM) ↔ EUDAT (IRODS/Griffin technology with GridFTP extension)
- Data Transfer using
 - FTS3 (using SRM protocol on EGI side and pure GridFTP on EUDAT side)
 - Transfer using GlobusOnline.eu (using pure GridFTP protocols on both infrastructures. Data transfers between EGI and PRACE were not tested in practice, but the assumption was made that they are working, because PRACE is using pure Globus GridFTP protocol.

6 Objective 3: Integration of common operational services

European e-Infrastructures like EGI, EUDAT and PRACE are inherently distributed and share similar operational needs.

- Central operations can be provided for all e-Infrastructures, while the running operations services that are specific to an e-Infrastructure is delegated to the e-Infrastructures themselves. Benefit: cost reduction and no duplication.
- Operational tools of common interest are jointly developed as open source project. This allows innovation and improves sustainability.

High-level operational services. A large number of common operational high-level activities and services that are not specific to a given infrastructure, service, technology or service provisioning model can be potentially shared and provided jointly. High-level operational services include: the service repository certification (GOCDB), accounting, monitoring and service level reporting, security coordination, incident response, security monitoring tools, messaging, security policy development, integrated incident management etc.

Tools that are of common interest can develop into open source projects that are evolved to meet the needs of different e-Infrastructures.

The Appendix I-II-III shows in detail how tools can be shared by different e-Infrastructures (monitoring, GOCDB and Accounting).

e-Infrastructure operational services. Operational services that are specific to an e-Infrastructure are independently provided

7 Objective 4: Integration of user-facing services

All e-Infrastructures are engaged in supporting user communities to foster the adoption of their services, evolve them according to the user requirements and ensure innovation, and in training programmes. These services are currently provided with little or no coordination. As a consequence exploitation and sharing of expertise available in different communities are difficult, and user communities need to approach different e-Infrastructure to develop an integrated solution.

User communities would benefit from a single virtual Centre of Excellence for:

- User and Resource Provider technical consultancy
- User requirements' analysis
- Training
- Resource provisioning to allow access to different distributed resource pools: currently different agreements have to be established with different service providers for data, grid, cloud, etc.

Consultancy and training could be provided in collaboration to external experts from the Technology Provider's community and the user community side.

ESFRIs and user communities could be not only consumer but also provider of

- consultancy services for other communities, see for example the reference model skills developed in the ENVRI project;
- technical services, whenever VRE solutions developed in house can be adopted or customized by other user communities.

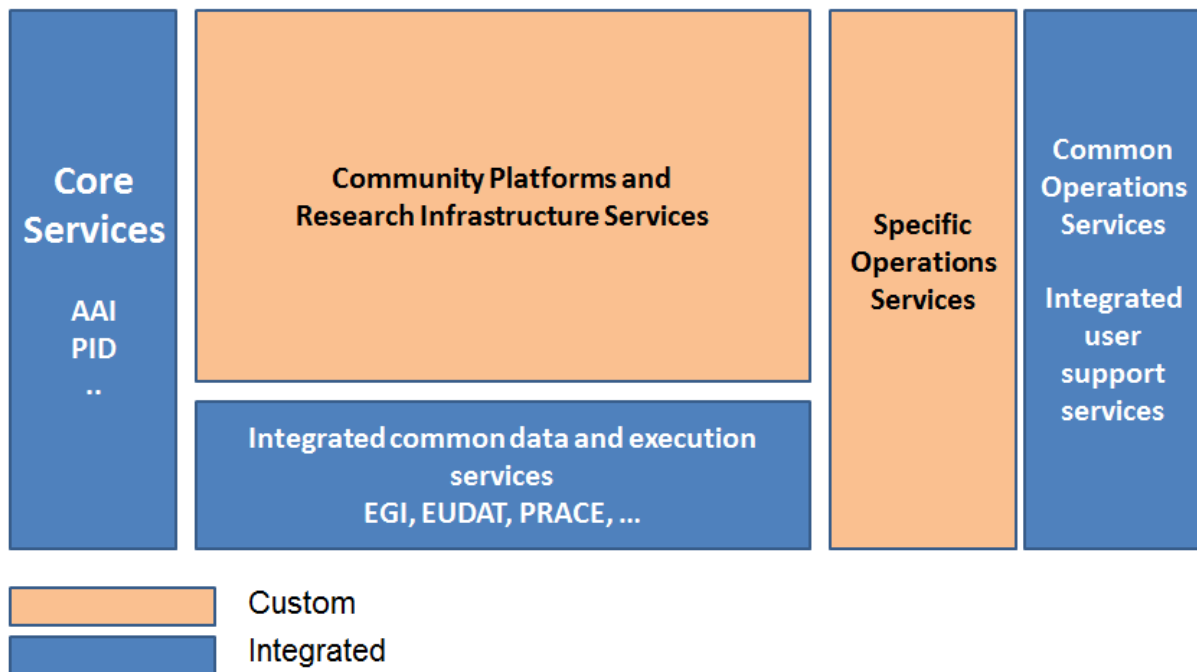
8 Objective 5: Business development and sustainability

Business development and migration to a model where users can use their purchase power to help sustainability of e-Infrastructures requires:

- An integrated framework for service level agreement negotiation
- Availability of a single broker interface for service provisioning
- Adoption of best practices for service management in a federated provisioning environment where assets like compute, storage, data and higher level services need to be provided jointly.
- Partnership with industrial customers

e-Infrastructures can share business models, share expertise in business development and develop an integrated approach for the establishment of a single European market place of services and resources for users.

9 Integration vision for discussion



10 Conclusions

This paper presents five different collaboration opportunities in the areas of service provisioning, operations, user services and sustainability.

Collaboration opportunities will be discussed by the e-Infrastructure representatives to establish a common integration roadmap to better serve user communities, and ensure coordination in the areas of testing, integration at a European level.

11 APPENDIX 1. Infrastructure Monitoring

GRNET, CNRS and SRCE are developing a new framework that is the evolution of the existing SAM framework. It addresses the evolving requirements of e-Infrastructures for testing and benchmarking their capabilities in terms of resilience and service continuity. The goal is to provide an end-to-end solution for service monitoring, alerting and reporting on large scale infrastructures.

The monitoring engine is fully compatible with the Nagios framework. Operators can use a wide range of existing monitoring probes or integrate their own. The monitoring engine will be able to auto-configure itself by retrieving information from external sources, such as the GOCDB, and can be used by operations management systems such as the Operations Portal. Furthermore, it supports sending notifications/alarms in case of service problems based on configurable criteria.

The reporting engine is a powerful analytics engine that can provide custom dashboard & reports for the status and the availability and reliability of the infrastructure. It supports down to hourly resolution of results and the generation of a wide range of charts that can provide valuable insights to both the managers and the operators of the infrastructures.

The development of the new framework is open and the core team behind it is soliciting requirements from possible users of the service. To this end, a requirements gathering group has been setup with representatives from e-infrastructures and scientific communities. The group holds open conference calls in regular intervals, in which there is an open discussion of the roadmap and the future goals.

12 APPENDIX 2. Service Registry (GOCDB)

A common information system to register resources and record topology data is essential to effectively manage a distributed computing infrastructure, especially across multiple projects, grids and clouds. Topology data such as sites, services, endpoints and downtimes should be defined once without duplication, and published in formats that are interoperable across different administrative domains such as EGI and EUDAT. This is essential to eliminate data inconsistencies and maintain the integrity of data on which other systems such as monitoring and accounting are built. For GOCDB, collaborative requirement analysis and joint development activities will ensure EGI and EUDAT continue to benefit from the advantages already recognised by a common information system. Importantly, this prevents similar features and requirements from being duplicated across projects. When divergence is necessary in certain circumstances, EGI and EUDAT have the opportunity to agree upon common abstraction points which prevents excessive tool-forking and helps facilitate per-project customisation. A recent example which successfully demonstrates the benefits of shared development includes the extensible GOCDBv5 'resource-scoping' feature, which allows resources to be tagged by one or more project scope-tags. This allows resources to be categorised into flexible project groupings using the same mechanism (i.e. 'EGI,' 'EUDAT' and 'Local' for no project affiliation). With H2020, the opportunity now exists to continue these development best-practices. Development areas that would both benefit both EGI and EUDAT include: a) further extension of the data model to define additional attributes/data, b) rendering of the GOCDB data in a standard format such as the recently published GLUE2 XML rendering from OGF (rather than with proprietary XML), c) additional REST endpoints to programmatically POST data as an alternative to manually inputting data via the Web portal, d) integration of additional authentication mechanisms, e) new abstractions to integrate project-specific business rules and roles, potentially using a rules-engine, f) GUI enhancements for mobile devices.

13 APPENDIX 3. Accounting

The EGI APEL Core Services provide a single place for projects running on distributed resources to store accounting information on a variety of resources. Some of them are ripe for exploitation by EUDAT science users and are already adopted/integrated by external peer e-Infrastructures like Open Science Grid in the US.

- CPU: APEL currently holds data on 2.5×10^9 jobs from 350 sites from between 2004 and 2013. The data comes from a variety of e-infrastructures (EGI, OSG, NeIC, INFN) and middleware stacks (Globus, gLite, ARC, Unicore, QCG, DesktopGrids). Also in discussions with XSEDE and PRACE. A relevant usecase for EU-DAT sites would be those sites who run cpu resources to allow users to do processing close to their data. Publishing cpu accounting to APEL would allow the scientific projects to aggregate their cpu usage across EU-DAT sites and beyond to other e-infrastructures. (Given that many EU-DAT resource centres are also EGI ones, it is possible that this is partially happening already.)
- Storage: sensors to collect usage information from the DPM and dcache storage systems were developed in EMI. APEL is currently testing these and also collecting storage usage data from the Information Service. A useful collaboration with EU-DAT would be to extend the range of storage systems supported to include i-Rods and any others used by EU-DAT.
- Data Access: The storage accounting mentioned above is only for occupancy of storage by projects, not for who uses the data, how often and from where. We call this latter set of data 'data access accounting'. EU-DAT might be interested in this type of accounting for their sites and projects and work together on a project with APEL to specify and develop such accounting.
- Cloud: APEL is recording the usage of the EGI Federated Cloud Infrastructure at the VM instance level. The technology to do this might be of interest for re-use by EU-DAT. Another scenario like the cpu one above is if EU-DAT sites are providing cloud resources close to their data to support the user who has their own (or commodity, or software) images that they wish to use to process their data.