



EGI-InSPIRE

FEDERATED OPERATIONS SOLUTION

EU DELIVERABLE: D4.9

Document identifier:	EGI-InSPIRE-D4.9-v1.3.docx
Date:	24/10/2013
Activity:	SA1
Lead Partner:	EGI.eu
Document Status:	FINAL
Dissemination Level:	PUBLIC
Document Link:	https://documents.egi.eu/document/1967

Abstract

This document constitutes a ‘White paper’ describing the Federated Operations Solution that offers the technologies, processes and the organisation required to manage the operations of a heterogeneous infrastructure. It also proposes how to integrate resources from multiple independent providers with a lightweight central coordination.



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	T. Ferrari, M. Krakowian	EGI.eu/SA1	21-10-2013
Reviewed by	Moderator: D. Scardaci Reviewers: S. Andreozzi O. Appleton	EGI.eu/JRA1 EGI.eu/NA2 Emergence Tech Limited	04-10-2013 09-10-2013 04-10-2013
Approved by	AMB & PMB	5	24-10-2013

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	02-10-2013	v1.0	T. Ferrari/EGI.eu
2	15-10-2013	v1.1-1.2 Revisions	M. Krakowian/EGI.eu
3	15-10-2013	v1.3 Final version for PMB review	M. Krakowian/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

EGI provides a Solutions Portfolio that addresses the needs of different customer segments: national and international research collaborations, research infrastructures, individual researches and Resource Providers (the individual Resource Centres as well as federated providers like the National Grid Initiatives (NGIs) that are part of the EGI Collaboration).

Solutions are enabled by combining and delivering services from EGI.eu and NGIs Service Catalogues¹. The service components are provided by a number of different stakeholders of the EGI ecosystem: EGI.eu, NGIs, Resource Centres and Technology Providers through the EGI.eu coordination. The EGI.eu catalogue has been available online since May 2013, after an analysis of the EGI ecosystem to document all EGI activities and redefine them as services, according to their business value. The NGIs' catalogues are under development with pilot organisations following the experience and analysis performed by EGI.eu.

The EGI Solutions Portfolio includes:

- Federated cloud
- Federated operations
- High-throughput data analysis
- Community networks and support
- Community-driven innovation

This document describes the Federated Operations Solution, its value proposition, the targets groups addressed, the services that collectively contribute to the delivery of the solution, the strategic impact and the performance indicators.

The main goal of this document is to present EGI Federated Operations Solutions which can be reuse by other organizations to build the Federated Operations Solutions for their domain.

¹ <http://www.egi.eu/services/catalogue/>



TABLE OF CONTENTS

1	INTRODUCTION	6
2	TARGET GROUPS	7
3	THE CHALLENGES	8
4	FEDERATED OPERATIONS	9
4.1	Service sharing	10
4.2	Federated Service Management	12
4.3	Standards, protocols and interfaces	12
4.4	Policies and Procedures	13
5	VALUE PROPOSITION	14
6	IMPLEMENTATION	14
6.1	Operations Coordination	14
6.2	Technology Coordination	14
6.3	Security Coordination	15
6.4	Federated Operations	15
6.5	Helpdesk Support	16
6.6	Specialized Consultancy	16
6.7	Technical Consultancy and Support	16
7	STRATEGIC IMPACT	17
8	MEASURING SUCCESS	17
9	SUMMARY AND OUTLOOK	18
10	CONCLUSION	19
11	REFERENCES	19



1 INTRODUCTION

EGI provides a Solutions Portfolio that addresses the needs of different customer segments: national and international research collaborations, individual researchers, Research Infrastructures and Resource Providers – individual Resource Centres as well as federated providers like the National Grid Initiatives (NGIs) that are part of the EGI Collaboration.

Solutions are enabled by the delivery of services of the EGI.eu Service Catalogue [CAT], which was published on-line in May 2013. These are provided by a number of different stakeholders of the EGI ecosystem: EGI.eu, NGIs, Resource Centres and Technology Providers through the EGI.eu coordination. This catalogue defines all EGI activities and structures them as a service, according to their business value. However, it is important to stress that these solutions are not created from nothing for this purpose. They are a re-description, alteration and evolution of activities that EGI (both EGI.eu and EGI Inspire) carried out already, and indeed activities dating back to the previous EGEE and DataGrid project series.

The current EGI Solutions Portfolio includes:

- Federated cloud
- Federated operations
- High-throughput data analysis
- Community networks and support
- Community-driven innovation

This document provides a high-level description of the Federated Operations Solution that is intended for distributed Research Infrastructures as well as distributed e-Infrastructures that adopt a federated governance model.

The e-IRG white paper 2013 [eIRG] introduced the concept of “e-Infrastructure Commons” as a set of integrated services provided through the interoperability of e-Infrastructures with the aim to remove the existing technical, functional, geographical, institutional and political barriers. According to the description depicted in the document, the EGI Federated Operations Solution is able to satisfy various e-IRG technical recommendations to create the “e-Infrastructure Commons” and EGI.eu aim is to continue the evolution of the EGI Federated Operations Solutions to cover more e-IRG, and generally speaking EC, requirements.

The EGI Federated Operations Solution is the output of a transformation process, taken place during EGI-InSPIRE, that modified the old Grid operations used during the EGEE projects in a federated solution, independent from the underlying technologies, and able to serve different Research and Resource Infrastructures.

Federated e-Infrastructures are a collection of at least partially *autonomous* entities that act together in some manner to provide a consumer with services. This federation includes a central *federator*, whose role varies widely depending on the type of federation considered, from strong central control through to weak coordination or even internal support for the federation and its members.

The Federator is a body that provides some value-added services that somehow relate to the whole federation. These can range from very abstract or high level services such as advising customers on their needs, through fully integrated services offering a single point of contact to customers and concealing the underlying federation.

The European Grid Infrastructure is an example of federated system (Figure 1), where individual Research Centres are autonomously operated nationally by their National Grid Initiatives (NGIs).

Each NGI is fully responsible for its individual service provision and management as they are autonomous providers with their own (local) customers. However, the provision of common, federated services requires the adoption of common policies, procedures, strategies, service management and technical roadmaps. However, the adoption of common policies and procedures, strategies, technical roadmaps and the provisioning of central common services ensure integration of the national infrastructures. EGI.eu² is the central organization that plays the role of Federator.

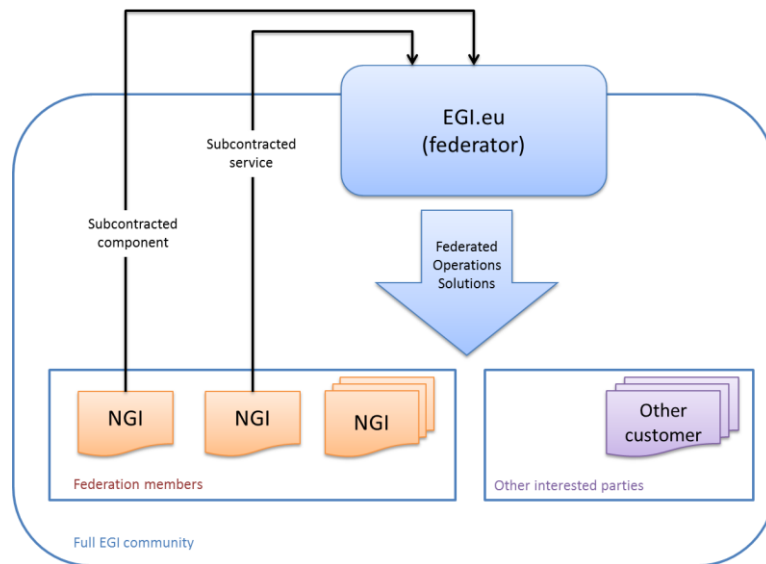


Figure 2 Visualization of the relationship in the EGI federation

This document is structured as follows: Section 1 is this Introduction; Section 2 defines the target groups addressed by Federated operations.; Section 3 defines the problems addressed; Section 4 outlines the proposed solution; Section 5 and 6 present the value proposition and the services of the EGI.eu portfolio that are used to deliver the solution; Section 7 and 8 define the expected strategic impact and the performance indicators; Section 9 provides information about how the solution is being adopted to federate resources and finally Section 10 draws up the conclusion.

2 TARGET GROUPS

The Federated Operations Solution addresses the needs of distributed *e-Infrastructures* and general-purpose *e-Infrastructures*.

- **Research Infrastructures** [RI] refers to *facilities, resources* and related *services* used by the scientific community to conduct top-level research in their respective fields, ranging from social sciences to astronomy, genomics to nanotechnologies. These offer unique research services to users from different countries and help to shape scientific communities.

Examples include singular large-scale research installations, collections, special habitats, libraries, databases, biological archives, clean rooms, integrated arrays of small research installations, data infrastructures, research vessels, satellite and aircraft observation facilities, coastal observatories, telescopes, synchrotrons and accelerators, etc.

The Federated Operations Solution addresses the needs of distributed Research Infrastructures that encompass a network of physically distributed resources.

² <http://www.egi.eu/about/EGI.eu/>

- **e-Infrastructures** [EI] refers to the facilities and the related services used by multiple scientific communities which deliver core multi-disciplinary capabilities. Examples are high-capacity/high speed communication networks and highly distributed capacity and capability computing facilities.

What 'needs' do these target groups have in common?

- They need to coordinate the operations of distributed facilities or resources that span across organisations and countries.
- The coordinated operations services include accounting for usage, monitoring, support and incident management, infrastructure oversight, security operations etc.
- They need to offer predictable quality of service as a federation to be accessed by one or more research communities.

Evolving a solution portfolio

One difficulty faced in serving these groups in the context of a new solution portfolio is that EGI must evolve the current services and activities into more sustainable solutions without turning off the services or losing current users or customers. This is a difficult situation as user communities have experiences and consequent expectations about EGI services, not all of which may hold true with the new solutions. In some aspects EGI may add elements of the solutions that were not previously available, and here the new features must be promoted. In other areas it may ultimately make sense to withdraw features, especially in a time of lower budgets when not all features make economic sense (though low uptake versus cost, or thorough not being mapped to the realistic and sustainable strategic priorities of EGI). This change must be managed carefully but will inevitably perturb relationships with users to some extent, despite being a solid approach for longer term success.

3 THE CHALLENGES

Research Infrastructures and e-Infrastructures share common problems:

- **Lack of integration:** Science is multi-disciplinary in nature and requires the development and deployment of integrated, secure, permanent service-driven yet heterogeneous e-Infrastructures. However, the development of ad-hoc solutions to common operational needs such as monitoring, accounting, authentication and authorization and technical support from multiple providers is a barrier towards greater integration. The problem that is often faced is lack of interoperation and duplication of services; introducing harmonization at a late stage requires substantial effort and coordination with other organizations.
- **Lack of expertise and effort:** The delivery of production-quality services according to service management best practices and standards, and the development of an operations infrastructure require a substantial amount of expertise and effort. Small resource providers that are federated in an international infrastructure often do not retain enough expertise and effort.
- **Beta-testing in production:** It is the nature of research infrastructures and research-targeted resource infrastructures that they are developed incrementally in collaboration with user communities. As a result, instead of preparing a full service which is then released in its entirety as a finished product, these communities tend to develop as they operate, in effect beta-testing in production. This changes the dynamic of the relationship between the service provider and customer.

Identified common problems cause the following pains:

- **Inefficient or unpredictable service provisioning:** Services with similar functionality are developed and deployed by different e-Infrastructures what increase spending by increasing maintenance and effort cost.
- **Unpredictable service quality:** Due to diverse distribution of expertise and effort among Resource Providers capability to increase the quality of their service delivery is limited and may essentially vary.
- **Limited functionality or serious bugs in production:** Users see early versions of products with limited functionality or serious bugs that may make them less confident in the solution than if they only saw polished, finished products.

4 FEDERATED OPERATIONS

The Federated Operations Solution of EGI provides the following services for the standard operation of heterogeneous infrastructures from multiple independent resource providers with lightweight central coordination (Figure 3).

- **Technologies.** IT services which enable platform integrators to build different higher-level infrastructures (virtual research environments).
- **Processes.** Lightweight coordination is necessary to ensure the adoption of consistent policies, procedures and service management best practices.
- **Information.** Operations of a federated infrastructure require information about the distributed resources and services to be seamlessly available to customers, service providers and decision-making stakeholders.
- **Activities.** The evolution of federated operations is enabled through human services such as support, training, the availability of collaborative tools that facilitate the good functioning of working groups, task forces, etc.

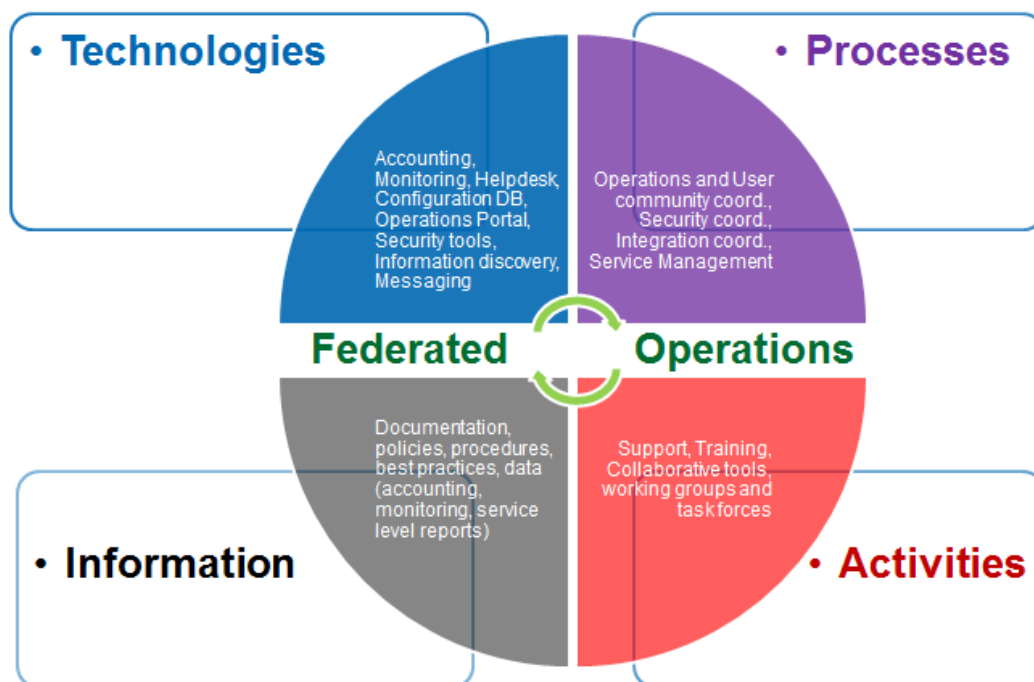


Figure 3. Assets delivered by the Federated Operations Solution of EGI

In EGI, the operation of heterogeneous distributed infrastructures is enabled through a set of operational interfaces that are service agnostic, and allow the central gathering of information for service registration, monitoring, accounting and support. The definition of these interfaces is standard-based when possible and evolves according to the needs of the infrastructures operated. These operational interfaces allow the adoption of custom tools and procedures for the operation of the local infrastructure, while being part of an international federation.

4.1 Service sharing

Within the EGI Federated Operations solution, duplication of operational effort and services is avoided by sharing a set of central services, activities and information. This section provides examples of how standard operations services can be provided centrally with EGI Federated Operations to ensure integration and avoid duplication.

- Customer and infrastructure operator's support.** The user support infrastructure in use within the EGI Helpdesk is distributed consisting of various topical and regional helpdesk systems that are linked together through a central integration platform, the GGUS helpdesk [GGUS]. This central helpdesk can be integrated with infrastructure-specific helpdesks and with Technology Provider helpdesks, and support teams in different thematic areas can be enabled. GGUS provides formalized communication between the submitter of the incident record and all partners involved in user support by providing an interface to which all other tools can connect and enabling central tracking of a problem, independent of the origin of the problem and the tool in which the work on the problem is done. The interlinking of all ticket systems in single place within EGI enables the passing of incident records from one system to another in a way that is transparent to the user (Figure 4). By exposing agreed interfaces, a tree hierarchy of interworking helpdesk systems operated by different infrastructures can be implemented allowing for transparently exchanging incident records. It also enables communication and ticket assignment between experts from different areas (e.g. middleware experts and application experts) while at the same time allowing them to work with tools they are used to, as well as the assignment of incidents to the affected infrastructure.

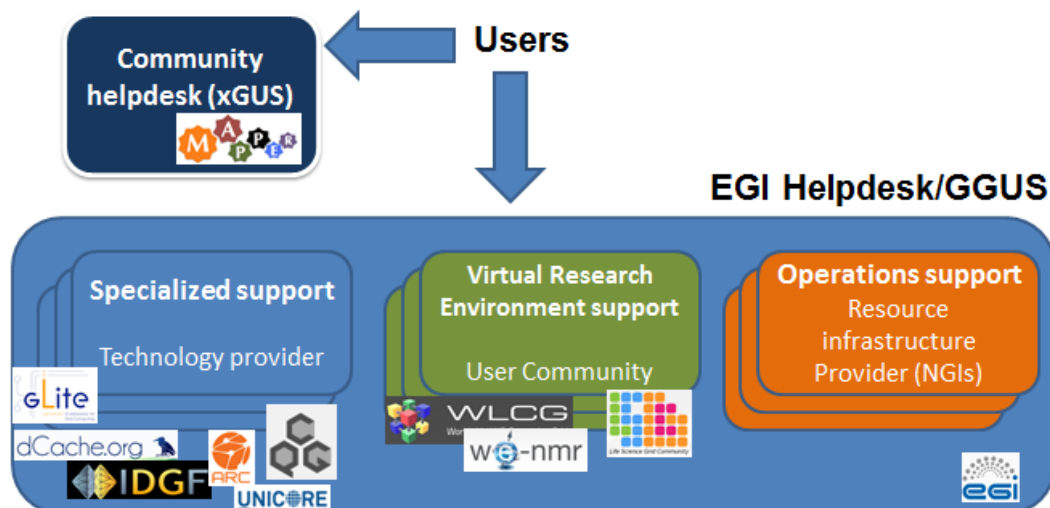


Figure 4. The EGI integrated helpdesk system. The service can provide integration between helpdesk systems of different Research Infrastructures (WeNMR, WLCG, etc.), e-Infrastructures (like PRACE, EUDAT etc.), and different technology providers offering 3rd level support.

- Accounting.** The EGI Accounting Infrastructure currently collects CPUs accounting usage records³ from Resource Centres and/or infrastructures and summarizes the data by Resource Centre, date (especially by month), VO, and user. This summary data can be displayed in a central Accounting Portal⁴ by dynamic queries on the parameters above at any level of the hierarchical tree structure that defines EGI and its partner grids. The central accounting portal provides a single point of access to infrastructures users and operators that are interested in an integrated view, which is particularly valuable to user communities that make use of services provided by different infrastructures. For example grouping accounting data from distributed infrastructures supporting international collaborations. The core EGI Accounting Infrastructure is based on APEL [APEL]. Other accounting systems in use by federated infrastructures can interface to APEL by publishing data in the central repository. This is possible by either publishing directly or by exporting a subset of the relevant data to the central EGI accounting repository.

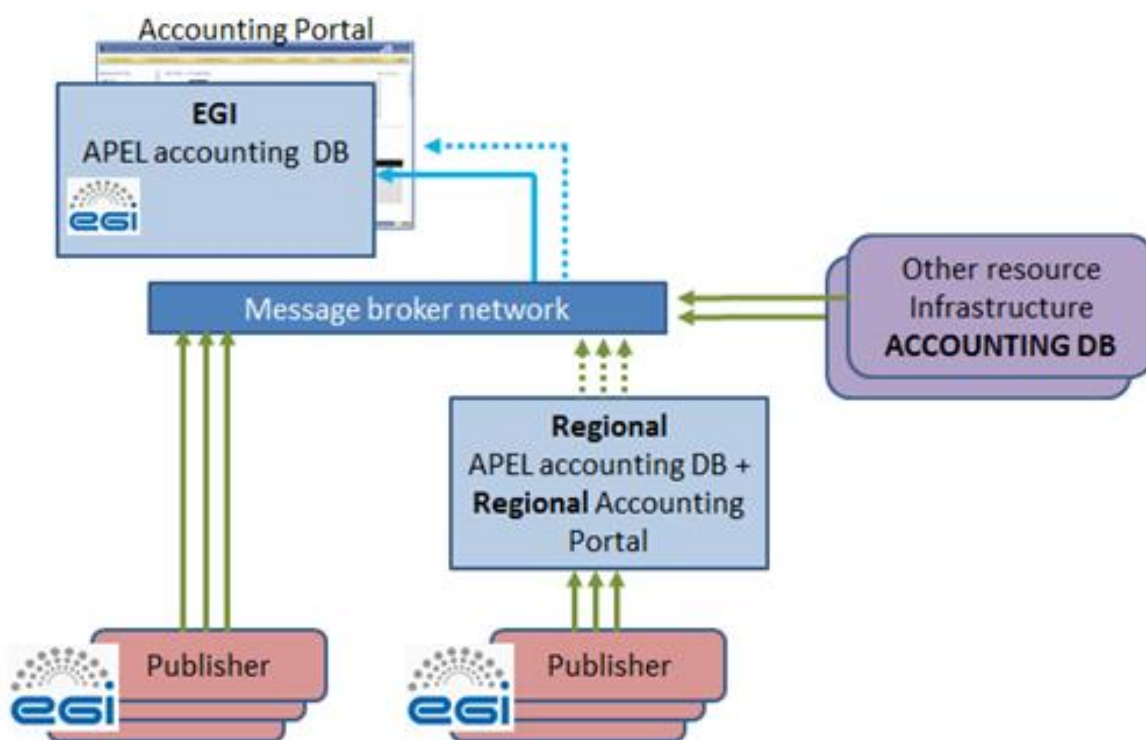


Figure 5. The EGI accounting infrastructure is a distributed system with central aggregation, where Resource Centres and federated infrastructures, like the National Grid Infrastructures, publish their Usage Records centrally by means of messaging. This infrastructure can be easily extended by plugging more infrastructures. Both the EGI central accounting repository and the Accounting Portal are consumers of accounting records exchanged between trusted publishers through the message bus.

³ The accounting system is also being extended to account for use of other types of resources: storage space and cloud. For more information about the status of EGI accounting integration please see [MS427].

⁴ <http://accounting.egi.eu/egi.php>

- **Central Site Configuration Database⁵** (GOCDB) [GOCDB] contains general static information about the Resource Centres (also known as “sites”) participating to one or more infrastructures, and most importantly, it provides topology information about them by defining the respective list of service end-points. GOCDB allows sites to store, maintain and view the topology of the production infrastructure and the basic site level information about the respective resources within it. GOCDB can provide information for multiple target infrastructures and is independent from the services being registered

4.2 Federated Service Management

Federated service provisioning is facilitated through the adoption of service management best practices, which are to establish common roles, document processes and agree on common service levels. The EGI solution provided aims at satisfying the requirements of the FitSM standard family [FSM]. The FitSM Standard requirements are aimed at harmonizing or improving service management across federated computing infrastructures. By this we mean situations where rather than a service provider having hierarchical conventional agreements with suppliers and customers, services are provided by a looser collaboration of organisations acting as a federation.

FitSM is a lightweight standard family aimed at facilitating service management in federated IT service provision. FitSM family is produced by the FedSM project⁶, an initiative co-funded by the European Commission Seventh framework Programme to improve service management in a selected set of federated ICT infrastructures and bring experience from this improvement to a broad community of federated communities.

4.3 Standards, protocols and interfaces

Interoperation is facilitated by the adoption of standard interfaces and protocols where possible.

EGI accounting usage record formats are based on Open Grid Forum⁷ standards: Compute Accounting Record [CAR] and Storage Accounting Record [STAR].

- In order to collect accounting data in a secure way, Resource Centres publishing usage records need to be registered in a central service configuration database.
- Exchange of records is supported by messaging, through the EGI's ActiveMQ⁸ Message Bus. *ActiveMQ* is an open source (Apache 2.0 licensed) message broker which fully implements the Java Message Service 1.1 (JMS). The APEL repository will accept the records into a holding container from where they will be merged with the summaries from other RIs and the summary produced by APEL from the job records it has received.
- The central accounting database exports data to the Accounting Portal where they can be viewed. Some accounting data are confidential and access can be restricted depending on the role of the viewer.

⁵ <http://goc.egi.eu/>

⁶ Implementing service management in federated e-Infrastructures (<http://fedsm.eu/>)

⁷ <http://ogf.org/>

⁸ <http://activemq.apache.org/>



Information discovery is a capability that helps find the required resources that have been registered with it within the production infrastructure. The information collected about such resources is made available through well-known instances that provide the data to some logical collection, infrastructure wide, regional, site, domain, etc.

Clients to such service must be able to search, filter, and order the available information until their initial request is satisfied. The LDAPv3 (RFC 4530) protocol and search syntax is used to query information from the information discovery services and to encapsulate the information payload relating to the services being offered within the production infrastructure that is exchanged between instances.

- EGI adopts the GLUE specification as an information model for Grid entities described using the natural language and UML Class Diagrams. GLUE, as a conceptual model, is designed to be independent from the concrete data models adopted for its implementation. EGI is currently based on version 1.3 of the specification [GLUE1], and is progressively adopting version 2.0 [GLUE].
- The EGI GLUE 2.0 profile [PRO] defines the usage of the GLUE 2.0 schema in EGI. It extends the schema specification document with more detailed semantics for particular attributes, specifies conditions under which objects and attributes should and should not be published, and defines some additional information to be published. It also includes guidelines for validating the accuracy of the published information.

The distributed monitoring infrastructure of EGI called Service Availability Monitoring⁹ (SAM), is based on Nagios [NAG].

- To integrate new services with SAM, tests for these service types have to be developed to cover the relevant functionalities defined in the management interface of the product. The probes are subsequently integrated into the SAM Release.
- Integration requires probe's developers to provide naming and test configuration (e.g. probe parameters, execution frequency, timeout, etc). SAM currently supports many Nagios probes [PROL]. Additional information about Nagios probe development and integration can be found in [PDI].

The Grid Site Configuration Database¹⁰ (GOCDB) [GOCDB] contains general static information and topology about the Resource Centres (also known as "sites") participating in one or more infrastructures.

- GOCDB information is made available in various ways, among which it is also rendered in GLUE 2 XML and JSON formats. New Programmatic Interface methods were developed to export information following the GLUE 2 format.

4.4 Policies and Procedures

Policies are needed to govern the provision of a high-quality distributed-computing production-oriented infrastructure. EGI.eu, the coordinating body for this community, 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.

⁹ <https://wiki.egi.eu/wiki/SAM>

¹⁰ <http://goc.egi.eu/>



5 VALUE PROPOSITION

The EGI Federated Operations Solution enables cost efficient operations in a federated environment. Central coordination is only needed to ensure integration among the Resource Providers; the responsibility of operating the federated infrastructures is shared with the providers, who also retain full control. The Solution allows increasing predictability of service quality and decreasing expenses.

Problems are removed by providing:

- **For lack of integration:** core Infrastructure Platform based on standards, common interfaces and protocols, communication, planning and coordination.
- **For lack of expertise and effort:** federated service management best practices, cost-effective sharing of services (support, processes, policies, activities), community expertise & re-use of tools/output from public funded projects.
- **For beta-testing in production:** validation support (processes, policies, activities, expertise), re-use of tools/output from public funded projects.

6 IMPLEMENTATION

The Federated Operations Solution for EGI is built by composing a number of services from the EGI.eu Service Catalogue [CAT] that are either centrally provided or offered by the infrastructures that are part of the federation. Other organizations facing the same problems could reuse EGI.eu services to build the Federated Operations Solutions for their domain.

6.1 Operations Coordination

Operations Coordination synchronises operational activities across the infrastructure to ensure seamless integration of services and to minimise fragmentation across the federation.

EGI Operations binds the infrastructure across Europe by ensuring that all technical services are delivered at an agreed service level and that computing services and resources are both seamlessly integrated and evolving according to the needs of the user community. This service currently provides EGI.eu members and resource providers with the coordination and management required to support more than 30 distributed Operations Centres, including development of policies and procedures.

Supporting activities

- Operations management coordination
- Operations integration coordination
- Coordination of tool development
- Service level management
- Grid oversight
- Documentation
- Network monitoring and support

6.2 Technology Coordination

Technology Coordination ensures continuous technological innovation through sourcing of software components from technology providers to meet the current and emerging needs of both researchers and resource providers.

EGI partners with external technology providers to source the software deployed across the production infrastructure. This includes prioritising technology requirements from user communities, commissioning the new software, as well as assessing and releasing the new products. The EGI.eu Technology Coordination Board (TCB) manages this process on behalf of EGI.eu members and coordinates high-level technology evolution across the EGI in order to ensure continuous innovation.

Supporting activities

- Supplier management
- Technical road-mapping
- Software provisioning

6.3 Security Coordination

Security Coordination ensures a secure and stable infrastructure to mitigate threats, enhance services, and give users the protection and confidence they demand from a service.

A secure e-infrastructure federation is one of EGI's top priorities. Guaranteeing a secure working environment for researchers requires coordination at various levels to prevent and handle incidents. The Security Policy Group is responsible for developing security guidelines and procedures. The Software Vulnerability Group aims to eliminate vulnerabilities from deployed software and prevent new ones. The EGI Computer Security Incident Response Team coordinates operational security, focusing on incident response, security monitoring, security training and security drills. This service ensures individual actions do not endanger the usage of shared systems and services.

Supporting activities

- Security operations coordination
- Security policy coordination
- External security organisation liaison
- Security incident response
- Security training

6.4 Federated Operations

Federated Operations simplify the day-to-day operations of a federated heterogeneous infrastructure avoiding duplication of costs and providing re-usable tools.

Federated operations brings together the operational tools, processes and people necessary to guarantee standard operation of heterogeneous infrastructures from multiple independent providers, with lightweight central coordination. This includes, for example, the monitoring, accounting, configuration and other services required to federate service provision for access by multiple research communities. A federated environment is key to uniform service and enables cost-efficient operations, while allowing resource centres to retain responsibility of local operations.

Supporting activities

- Accounting portal and repository
- Catch-all grid services for small user communities
- Development of operations monitoring probes
- Grid Configuration Database (GOCDB)
- Incident management tool (EGI Helpdesk)
- Message broker network
- Metrics portal



- Operational tools and meta-service monitoring (Ops-Monitor)
- Operations portal
- Security monitoring tools
- Service Availability Monitoring (SAM)
- Tools and services for resource centre certification

6.5 Helpdesk Support

The EGI Helpdesk Support offers professional, reliable and efficient technical support to guarantee a well-run infrastructure with improved productivity and usability for the customers. It provides a single entry point for problem and incident management, and can be integrated with the local helpdesks of the infrastructures that are part of the federation. It provides a mechanism to support both the use of the core infrastructure platform, as well as the community platforms deployed by the users, and their Virtual Research Environments by hosting community-specific support units.

EGI.eu coordinates the distributed user and operations support activities provided by the individual infrastructures and Resource Centres to ensure that operational issues are properly handled at both the Resource Centre and NGI level. This includes a centralised helpdesk ticketing system with dedicated oversight and follow-up.

Supporting activities

- Incident management tool (EGI Helpdesk)
- Multi-level support
- Ticket oversight and follow-up
- Ticket triage and assignment
- Network support

6.6 Specialized Consultancy

Specialized Consultancy offers tailored technical and management advice to help partners and clients make the most out of e-Infrastructure technologies.

Building on more than ten years of experience in e-infrastructure management, EGI.eu offers consultancy services on both technical and non-technical issues ranging from governance and policy to technical integration and support. The consultancy services can also take the shape of tailored advice to a wide of variety of organisations and projects that have EGI.eu as a strategic partner.

Supporting activities

- Governance support and consultancy
- Project support/consultancy

6.7 Technical Consultancy and Support

Technical Consultancy and Support helps research communities take the first steps in working with the infrastructure by providing the best solutions for their requirements and get scientific applications up and running.

Consultancy and support are essential to reach out to new user communities while maintaining the current ones. This service helps user communities to understand how they can engage and make the most of EGI resources. Consultancy offers technical advice for the best solutions, helps to get scientific applications up and running and provides general support mechanisms. The EGI.eu User Community Board is also tasked to gather, analyse and feed user requirements into the Technical Coordination Board for prioritisation and implementation.

Supporting activities

- User community coordination
- Requirements gathering and analysis
- Technical solutions consultancy (e.g. science gateways; workflows)
- Coordination of application porting

7 STRATEGIC IMPACT

The EGI Federated Operations Solution contributes to the EGI strategy by implementing a *secure federated data-analysis capability for the European Research Area*. The solution allows the operation of EGI capabilities for high-throughput data analysis by also extending it to other Research and e-Infrastructures, including the EGI federated cloud infrastructure, which according to the EGI vision, will scale up to 10 M computing cores and 1 EB of storage by 2020.

The European Research Area (ERA) will need to support researchers from diverse scientific disciplines taking approaches to data analysis. The EGI Federated Operations allow these to work seamlessly together in a distributed multi-disciplinary research collaborations that cross national and intellectual borders to tackle society's grand challenges.

EGI Federated Operations contribute to the ERA by increasing the ability of Europe to produce 'excellent science', which delivers exploitable innovations and new growth, Europe's researchers will need easy to use integrated services that provide access to high capacity and high quality computing and storage resources, wherever the resources and the researcher are located.

8 MEASURING SUCCESS

Table 1 provides the performance measures used to understand and define the efficacy of a proposed solution and its uptake. This set of measurements will be expanded and revised in the next update of the EGI Balanced Scorecard. The full list of EGI strategic metrics is available in D1.13 (Quality Plan and Project Metrics).

Objectives	Description	Performance measures
Improve operational efficiency and effectiveness	Improve the reliability and the delivery of the operational infrastructure through improvements in the operational tools and associated processes.	<ul style="list-style-type: none"> • Number of EGI Global Services meeting published Operations Level Agreements (OLAs)
Promote the sharing and re-use of innovation	Improve the reuse of innovation developed within the EGI ecosystem elsewhere in the ecosystem across all stakeholders (e.g. resource centres, research communities)	<ul style="list-style-type: none"> • Number of agreements established with external research communities to use EGI's operational tools to monitor their deployed services in their infrastructures

Table 1. Performance Indicators from the EGI strategic metrics measuring the success of the Federated Operations Solution.

9 SUMMARY AND OUTLOOK

- Federated Operations are used to operate distributed resources provided across 55 countries and one European Intergovernmental Research Institute (CERN).

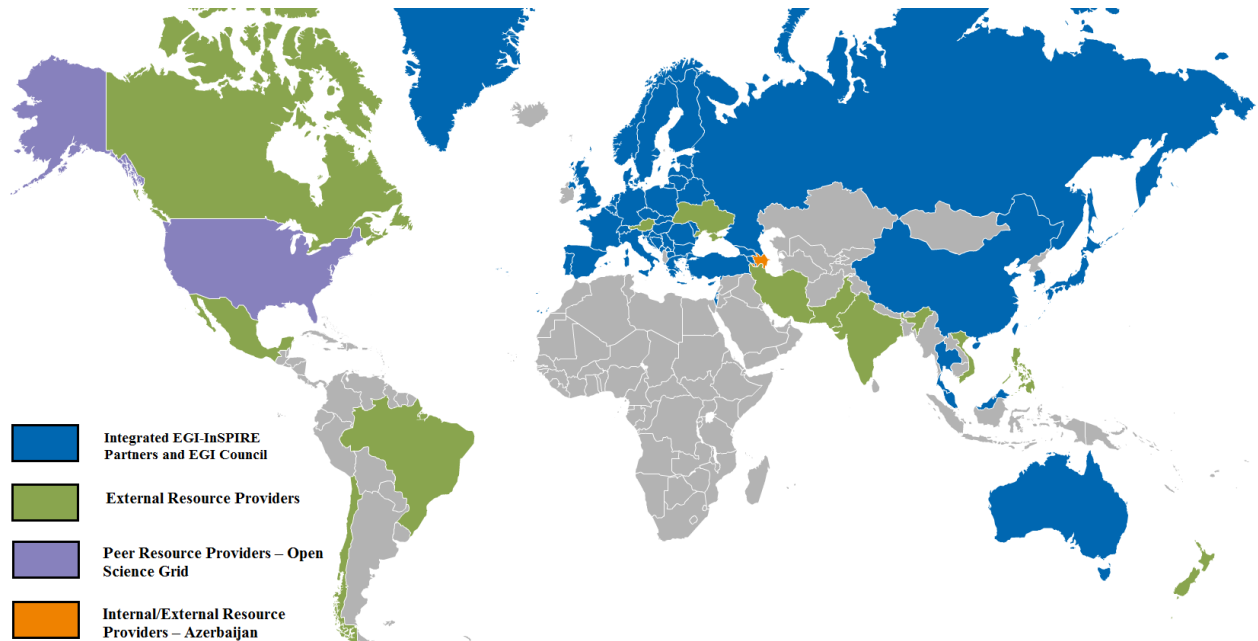


Figure 6 Resource Provider distribution in August 2013 (data source: GOCDB).

- EGI operates a fully integrated accounting infrastructure which services 38 distributed operations centres that oversee the resource provisioning by 340 Resource Centres distributed worldwide. It provides information about compute capacity utilization (local job, grid job and multi-threaded jobs) for both grid and desktop grid resources, and is being extended to collect usage statistics of storage and cloud resources. Accounting records are exchanged on-demand with peer infrastructures like the Open Science Grid¹¹ in the USA. Accounting.
- The EGI service registry (GOCDB) was adopted by EUDAT¹² to support operations, and EGI-InSPIRE supported the implementation of EUDAT requirements. EGI is currently responsible of the technical installation of the service. PRACE¹³ expressed interest in GOCDB testing. Collaboration with EUDAT is in progress about the evaluation of the EGI Service Availability Monitoring and its suitability to EUDAT deployment needs. EGI is actively collaborating with various ESFRI projects to investigate and demonstrate the adoption of the Federated Operations Solution.

¹¹ <http://opensciencegrid.org/>

¹² <http://www.eudat.eu/>

¹³ <http://www.prace-ri.eu/>



10 CONCLUSION

The adoption of EGI Federated Operations, deployed and developed within the EC project EGI-InSPIRE, allows service sharing, facilitates the adoption of federated service management best practices, policies and procedures, while relying on standards, protocols and interfaces where available. EGI can provide a solution for Federated Operations in a variety of configuration options, and can operate the solution in a software-as-a-service fashion for other infrastructures, or alternatively can service them through the existing production services of EGI.

Contact: **operations(at)egi.eu**
 EGI.eu
 Science Park 140, 1098 XG Amsterdam
 The Netherlands

11 REFERENCES

APEL	Roger, W., Jiang, M., Del Cano, C. et alt.; An APEL Tool Based CPU Usage Accounting Infrastructure for Large Scale Computing Grids, in the Proc. Of Int. Symposium on Grid Computing 2010, Taipei (TW), Mar 2010 http://event.twgrid.org/isgc2010/abstract/AnAPELToolBasedCPUUsageAccountingInfrastructureforLargeScaleComputingGrids.pdf
CAR	Usage Record - Format Recommendation, GFD.204, Open Grid Forum, 2013
CAT	EGI.eu Service Catalogue (http://www.egi.eu/services/catalogue/)
D1.13	Quality Plan and Project Metrics, EGI-InSPIRE Deliverable D1.13, June 2013 https://documents.egi.eu/document/1757
EGIP	Burke, S.; EGI Profile for the Use of GLUE 2.0 Information Schema, 2013 https://documents.egi.eu/document/1324
EI	Research Infrastructures programme e-Infrastructures activity http://cordis.europa.eu/fp7/ict/e-infrastructure/home_en.html
ERP	European Grid Initiative Resource Providers (http://www.egi.eu/infrastructure/resource-providers/)
eIRG	The e-IRG white paper 2013 http://www.e-irg.eu/images/stories/dissemination/e-irg_white_paper_2013_-_final_version.pdf
FSM	FitSM –Standard for lightweight service management in federated IT infrastructures http://fedsm.eu/fitsm
GOADB	Grid Configuration Database (GOADB) (https://wiki.egi.eu/wiki/GOADB)
GGUS	Global Grid User Support (GGUS) (https://wiki.egi.eu/wiki/GGUS)
GLUE	GLUE Specification v. 2.0, GFD-R-P.147, Open Grid Forum, 2009
GLUE1	Andreozzi, S.; Burke, S.; Field. L. et alt.; GLUE Schema Specification version 1.3



	http://redmine.ogf.org/dmsf/glue-wg?folder_id=19), Jan 2007
MS427	Integrating Resources into the EGI Production Infrastructure, EGI-InSPIRE Milestone MS427, August 2013 (https://documents.egi.eu/document/1894)
NAG	Nagios, the Industry Standard in IT Infrastructure Monitoring (http://www.nagios.org/)
PDI	Nagios Probe Development Guidelines (https://tomtools.cern.ch/confluence/display/SAMDOC/Probes+Development)
PRO	Burke, S.; EGI Profile for the Use of the GLUE 2.0 Information Schema (https://documents.egi.eu/document/1324)
PROL	SAM Probes: https://wiki.egi.eu/wiki/SAM#Tests_and_probes
RI	Infrastructures for Research and Innovation, European Commission (http://ec.europa.eu/research/infrastructures/index_en.cfm?pg=what)
STAR	StAR: Definition of a Storage Accounting Record, GFD.201, Open Grid Forum, 2013