

Demonstrating Excellent European Science on EGI's shared resources

Authors:

S. Andreozzi, T. Ferrari, D. Marinovic, S. Newhouse (EGI.eu)

Abstract:

The document provides a framework and recommendations for the application process, scientific review process and resource allocation process of NGI resources through EGI.eu (either pooled or paid).

Table of Contents:

1	For EGI Council Decision	2
2	Introduction	2
3	EGI Resource Allocation Process	3
3.1	<i>Eligibility of Customers</i>	3
3.2	<i>Project Type</i>	3
3.3	<i>Assessment of Scientific and Technical Merit</i>	3
4	EGI Resource Allocation Implementation	5
4.1	<i>Mechanism</i>	5
4.2	<i>Monitoring</i>	5
5	Impact of using EGI Resources	5
5.1	<i>Acknowledging Support</i>	6
5.2	<i>Open Access to Research Products</i>	6
6	Open Issues and Next Steps	6
7	References	6
	VT Scientific Publications Repository - Recommendations https://documents.egi.eu/document/1369	6

Copyright notice

This work by EGI.eu is licensed under a Creative Commons Attribution 3.0 Unported License (see a copy of the license at <http://creativecommons.org/licenses/by/3.0>). This license lets you remix, tweak, and build upon this work, and although your new works must acknowledge EGI.eu, you do not have to license your derivative works on the same terms. Reproductions or derivative works must be attributed by attaching the following reference to the copied elements: "Based on work by EGI.eu used with permission under a CC-BY 3.0 license (source work URL: <https://documents.egi.eu/document/1415>)".

1 For EGI Council Decision

The key points within this report that need the endorsement of the EGI Council are:

1. Establishment of a 'Scientific Review Committee' (SRC) to manage the review of resource requests that will be assigned to EGI's pooled resources (see Section 3).
2. Identification of resources from NGIs and Resource Centres (RCs) that can be pooled together and managed through the 'Scientific Review Committee' with no further review from the NGIs or Resource Centres (see Section 4).
3. Mandate that *all* VOs and *all* users within a VO acknowledge EGI and the NGIs resources supporting their work by making this condition part of EGI's Acceptable Use Policy (see Section 5.1).
4. Establishment of a strategic relationship with the OpenAIRE collaboration to record publications in open access repositories by EGI, Research Discipline, and NGIs (see Section 5.2).
5. VOs unable to demonstrate evidence of their publications within open access repositories will not have their resource allocations renewed/extended and may lose access to other EGI services.
6. Task EGI.eu Operations to investigate mechanisms by which resource allocations made by individual sites to a VO can be managed and monitored so that the offered resources are actually delivered.

2 Introduction

EGI supports world-class research through over 200 Virtual Organisations (VOs) enabled on its distributed infrastructure across over 50 countries. Access to these computing resources is delegated to the relevant national resource allocation mechanisms. Research groups currently can obtain access to EGI resources through: 1) peer reviewed grants to directly purchase equipment to be installed and offered to them; 2) peer reviewed allocation of leased access to nationally funded equipment (either with guaranteed share or for opportunistic usage).

Currently, the Resource Centres (RCs) retain the ultimate decision on which VOs should be supported while VO Managers decide which individuals are able to access the resources within the VO. Users are expected to restrict their activities to those defined in the Acceptable Use Policy of the VO. While this model has been shown to scale to over 350 RCs and over 200 VOs to support over 21,000 users, it has some underlying drawbacks:

1. **Difficult to demonstrate scientific impact of resource allocations:** There are no systematic mechanisms at a *European* level (and frequently none at a national level) to capture the scientific impact of this work across VOs. This is now impacting the ability of national and European coordination bodies to demonstrate and promote their value.
2. **Complex allocation process:** All resource allocation decisions take place at a national level in potentially fragmented and complex process that fails to demonstrate European value of EGI; researchers need to talk individually to the NGIs (or RCs) in order to achieve cross-country resource allocation and definition of service level targets.
3. **Lacking a framework for establishing and monitoring resource allocations:** Resources centres either have a number of primary users which have dedicated resources (e.g. the case for the largest user community – HEP) or provide opportunistic access to multiple user communities with some (according to local allocation rules) or no (if not a locally strategically important user community) minimum resource quota. In the latter case the responsiveness of EGI to a user community's workload will be dependent on the workload of other user communities.

This report aims at providing a thought experiment around identify the application, assessment and allocation process that we could implement to support peer-reviewed science on resources pooled from the NGIs/EIROs. This resource allocation mechanism would be in addition to the current mechanisms used within EGI. This report provides a number of policy recommendations and technical actions to tackle these issues.

3 EGI Resource Allocation Process

The lifecycle of a *resource allocation* is bound to the concept of *project*. A project has a well-defined *customer* who initiates the request, a scope and may be assigned with one or more resource allocations during the project's lifetime. Within EGI the customer can be a community defined by a common purpose (the project) and a set of associated individuals *consumers* represented by a VO Manager and by VOMS. During the lifetime of the project, the allocated resources attached to a project can expire, be renewed, modified or added to depending on the needs of the consumers.

These terms are defined by:

- **Project:** a group activity with a defined beginning and end time, with allocated resources to achieve a well-define result.
- **Service Provider:** an organisation supplying services to one or more consumers. In our scenario, we distinguish two main types of service providers: 1) **Resource Provider:** an organisation offering access to ICT resources through service abstractions (e.g., computing power, storage); 2) **Broker:** an organisation arranging transactions and agreements between a customer and one or more resource providers (e.g. EGI.eu at a European level, and NGIs at a national level).
- **Customer:** the person or entity that negotiates and commissions the level of services with a service provider and may pay, doing so on behalf of a number of consumers (users)
- **Consumer:** the person actually using the service (user)

The following policies are defined relating to the actors:

3.1 Eligibility of Customers

The customer is responsible for the accuracy of the request and for the management of the ensuing allocation by its consumers. There are different categories of customers eligible for requesting EGI resources:

1. Researcher or educator affiliated to European academic or non-profit research institutions.
2. Employers of museums, observatories, libraries, research laboratories, professional societies and similar organizations in Europe that are directly associated with scientific research.
3. Employers from European for profit organisations (especially SMEs) with scientific, research or backgrounds who wish to use EGI resources for pre-commercial research.
4. Scientists, engineers or educators located within the Europe may be eligible for support, even if the individual is not employed by or affiliated with an organization provided that the project is sufficiently meritorious and otherwise complies with the conditions of any allocation request and the requestor has demonstrated the capability to carry out the project.

Researchers not based in Europe are not eligible for resource allocation requests, however they are eligible to participate in projects with approved resource allocation requests as consumers.

3.2 Project Type

Each project request must be categorised accordingly to one of the following areas:

- **Start-up project:** It targets new users planning to submit more substantial requests for EGI resources in the future as well as users who have modest computational requirements that, nevertheless, cannot be satisfied by their local resources. Resources allocated to such projects will be restricted and enforced by monitoring usage.
- **Research project:** It targets individual or groups that have progressed beyond the start-up phase and are conducting "production" usage of EGI resources in pursuit of their research objectives.

Other project types may be defined as the project request system develops.

3.3 Assessment of Scientific and Technical Merit

The scientific and technical merit of the work that is to be supported by the resources being requested by the project is assessed by the Scientific Review Committee (SRC). The composition of the SRC consists of scientific advisors affiliated to members of EGI.eu. A short list will be selected by the EGI.eu Executive Board for adoption by the EGI Council. The SRC review process will be:

1. The SRC will review the scientific merit of the project that will use the requested resource allocations:
 - a. If this is a new resource request not affiliated with an existing peer reviewed activity (e.g. and FP7 or H2020 research project, research infrastructure, etc.) at least two SRC reviewers will be appointed to assess the scientific merit if the project.
 - b. If this is a new resource request affiliated with an existing peer reviewed activity the scientific review will consist of just ensuring there is a link between the activity and the project request. This may be taken directly by the SRC Chair.
 - c. If this is a request to extend or renew an existing resource allocation, then it will rely on the basis of the peer-reviewed publications and their scientific impact recorded in EGI's Scientific Publications repository (see Section 5.2).
2. The SRC will review the technical merit of the approach being proposed. To assist the SRC in making this assessment, experts from the NGIs may be requested to provide their advice. This will consider:
 - a. For Compute Resources: Is it clear from the description and methodology that the choice of applications, methods, algorithms and techniques to be employed will accomplish the stated scientific objectives?
 - b. For Storage Requests: Is it clear from the description and methodology that the projected data usage, access methods, algorithms and techniques to be employed will accomplish the stated scientific objectives?

All resources requests should focus on the efficient use of resources and in accordance with the recommended user policies.

3. The SRC decision-making process is based on consensus and will be determined by the responses from the rest of the SRC to the reviews assessing the scientific and technical merit of the proposed project (see stage 1).
4. The requestors may appeal the decisions of the SRC to the EGI.eu Executive Board.
5. The Director (or delegate) will be an *ex-officio* member of the SRC and shall
 - a. Assign EGI.eu staff to support the work of the SRC and its Chair.
 - b. In conjunction with the Chair will report on the work of the SRC at regular intervals (at least annually) to the EGI.eu Executive Board and EGI Council.
 - c. In conjunction with the Chair monitor the review process to ensure that it is being conducted in an open and transparent manner that is free from potential conflicts of interest.
 - d. Ensure the confidentiality and anonymity of the review process.
 - e. Report to the SRC and the EGI Council the usage made by the approved projects and their individual resource allocations within the overall resources committed to the NGI to this activity.

The SRC and its reviewers shall evaluate the scientific and technical merits of the proposal based on the following information:

- The project must have more three or more identified consumers from organisations meeting the eligibility criteria defined in Section 3.1.
- The resource request will state the scientific impact of the research to be conducted.
- The technical approach that will be used to conduct the research on EGI resources will be described.
- Existing peer reviewed supporting grants (European or National) wholly or partially supporting the research should be listed.
- For extensions/renewals resource requests the SRC will also consider the progress made using prior allocations, including the peer-reviewed publications recorded in the designated EGI's Scientific Publications Repository.
- The alignment of the project to the priorities defined by the EGI Council for EGI resources.
- A resource allocation may be of any duration. However each year the publications produced by the project and registered in the open access publications repository will be reviewed. If no publication activity is found then the project's use of EGI resources may be suspended.

A project does not need to have the direct or indirect support of a funding agency to request resources, however a credible plan needs to be presented as to how the requested resources will be used to achieve the desired scientific goals.

4 EGI Resource Allocation Implementation

4.1 Mechanism

EGI.eu does not own resources to support the peer reviewed allocation process defined in Section 3. According to a recent survey on resource ownership and current allocation policies [R1] conducted among NGIs, in most of the countries resources are owned by RCs. There are three possible scenarios for resource ownership:

- **RC ownership:** Resources are owned by RCs and by the user communities that they support locally. Access by other user communities to these resources may be allowed depending on local policies of the RC. The NGI may provide coordination for the allocation of resources to user communities that do not own resources directly within the limits set by the local RC.
- **NGI ownership:** In a small number of cases, resources are directly owned by the NGI; these are typically supported through national funding channels which manage resource allocation.
- **NGI and RC ownership:** In several NGIs, resources are owned by both the NGI and RCs affiliated to the NGI. The NGI in some case offers resources to emerging user communities according to an opportunistic usage model.

There is no point establishing a centrally coordinated resource request process if the RC is going to impose further restrictions, reviews or delays on the requested allocation.

Therefore it is proposed that:

1. VOs are established on resources pledged by NGIs or individual RCs as part of this program that will be managed centrally by EGI.eu as a result of the scientific review.
2. Pledges made by an NGI or an individual RC will be reviewed annually and the usage of these pledged resources will be reported.
3. Approved projects and their users will be allocated to these VOs centrally through VOMS and VOMS groups.

4.2 Monitoring

While the centrally managed VOs introduced previously will provide access to resources at participating RCs, there is no guarantee that if the VO is only provided with opportunistic access with no minimum quota that it will be able to achieve any useful work. This is no different from the model currently provided within EGI for large diverse VOs such as BioMed. It is therefore proposed that:

1. RC's and NGIs acknowledged as contributing resources to this model must during its initial phases allocate non-trivial resources and monitor the usage (job volume & queuing time) to ensure a satisfactory experience is delivered to the user.
2. Consumption of resources at different RCs by a project will be reported regularly by both RC and by project to the EGI Council and the broader community.
3. EGI.eu Operations will investigate mechanisms by which resource allocations to a VO can be defined by a RC and monitored for delivery by both the RC and EGI.eu. This would provide a mechanism in the future for capacity planning around central allocations, but also benefit other VOs operating within EGI in the long-term.

5 Impact of using EGI Resources

Measuring the impact of EGI resources within the scientific community is critical for future European and National investment in EGI and the NGIs activities. The EGI Scientific Publications Repository Virtual Team recently made a number of recommendations in this area [R3] that are integrated into this broader initiative.

5.1 Acknowledging Support

Acknowledgment of the support provided by EGI and its cooperating NGIs must appear in the publication of any material (whether copyrighted or not) that is based on work supported wholly or partially using EGI resources. Specific acknowledgement text will be provided. The obligation to cite the support provided by EGI is mandatory and will become part of EGI's Acceptable Use Policy which by definition will impact all users in all VOs. Only customers who acknowledge usage of EGI in their publications deposited in open access repositories affiliated with OpenAIRE will be eligible to extensions of any allocated resources.

5.2 Open Access to Research Products

EU Open Access policy requires that scientific results (publications and data) are publicly accessible. EGI should leverage this requirement in order to help assess the scientific impact of its work. Through collaboration with OpenAIRE EGI.eu will work on a customised version of the OpenAIRE portal that will provide:

- A view of all publications connected to EGI
- Allow researchers by only providing the DOI publications categorised by 1) funding project, 2) infrastructure (default: EGI), 3) scientific discipline(s) and 4) virtual organisation(s) to which the publication is related to;
- Automatically harvest publications from repositories or publishers connected to OpenAIRE by adding text-mining capabilities searching for EGI citations for later verification.

6 Open Issues and Next Steps

A number of open issues still need to be resolved around this work

1. Must all customers (those submitting resource requests) originate in Europe (this does not mean that they have to be European) but consumers can be from anywhere in the world?
2. Are extensions needed to the Accounting Portal to view the usage made by the 'EGI Managed Resources' VO by RC and by VO group?

Following on from the endorsement of the EGI Council of this report, the next steps are to:

1. Establish an EGI-Managed Resources VO on cooperating resources where access will be determined by the Scientific Review Committee.
2. Devise and maintain a consistent classification of scientific disciplines that can be used across EGI's various affiliated services.
3. Task the EGI.eu Executive Board to develop detailed terms of reference for the Scientific Review Council and to propose a list of members for adoption by the EGI Council.
4. Task EGI.eu Operations to develop through the OMB detailed procedures around this policy.
5. Inform the VOs and their members about the changes in the AUP and the obligation that now exists to cite EGI and deposit publications within OpenAIRE and its affiliated repositories.
6. Task the EGI.eu Strategy and Policy Team to establish a long-term collaboration plan with OpenAIRE and collaborate on the creation of the customised Scientific Publications Repository to meet the need of the EGI community

7 References

R 1	Resource allocation survey taken by the EGI.eu Operation Team in September 2012 https://wiki.egi.eu/wiki/Operations/Resource_Allocation
R 2	The EC Communication Towards better access to scientific information: Boosting the benefits of public investments in research http://ec.europa.eu/research/science-society/document_library/pdf_06/era-communication-towards-better-access-to-scientific-information_en.pdf
R3	VT Scientific Publications Repository - Recommendations https://documents.egi.eu/document/1369
R 4	XSEDE Allocation Policies https://www.xsede.org/web/guest/allocation-policies