



**HAL**  
open science

## Resource Allocation for the French National Grid Initiative

Gilles Mathieu, Hélène Cordier

► **To cite this version:**

Gilles Mathieu, Hélène Cordier. Resource Allocation for the French National Grid Initiative. Lecture Notes in Computer Science, 2012, 7156, pp.55-63. 10.1007/978-3-642-29740-3\_8 . hal-00714976

**HAL Id: hal-00714976**

**<https://hal.science/hal-00714976>**

Submitted on 6 Jul 2012

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Resource Allocation for the French National Grid Initiative

Gilles Mathieu, H el ene Cordier

IN2P3/CNRS Computing Centre  
43 bd du 11 Novembre 1918, 69622 Villeurbanne Cedex, France  
{gilles.mathieu, helene.cordier}@in2p3.fr

**Abstract.** Distribution of resources between different communities in production grids is the combined result of needs and policies: where the users' needs dictate what is required, resource providers' policies define how much is offered and how it is offered. From a provider point of view, getting a comprehensive and fair understanding of resources distribution is then a key element for the establishment of any scientific policy, and a prerequisite for delivering a high quality of service to users.

The resource allocation model which is currently applied within most national grid initiatives (NGIs) was designed with the needs of the EGEE (Enabling Grids for E-science) projects and should now be revised: NGIs now especially need to assess how resources and services are delivered to their national community, and expose the return on investment for resources delivered to international communities.

The French NGI "France Grilles" is currently investigating down this route, trying to define key principles for a national resource allocation strategy that would answer this concern while allowing for the proper definition of service level agreements (SLA) between users, providers and the NGI itself.

After looking for clear definitions of who are the communities we are dealing with, we propose to look at how resource allocation is done in other environments such as high performance computing (HPC) and the concepts we could possibly reuse from there while keeping the specificities of the Grid. We then review different use-cases and scenarios before concluding on a proposal which, together with open questions, could constitute a base for a resource allocation strategy for the French national grid.

## 1 Context and definitions

### 1.1 Context of current work

The EGI-Inspire [1] project started in May 2010, as a continuation of around 6 years of EGEE projects [2]. In this context, the French National Grid Initiative "France Grilles" [3] as emerged as EGI's partner for federating and operating grid resources in France. Within EGI operational context [4], allocation of

resources to grid user has changed scope, since national grids are now privileged interlocutors and interfaces between users and providers. This work is a preliminary reflection on the topic of resource allocation, and a possible basis for establishing policies and procedures in the medium term specific to France’s context and based on international collaboration.

### **1.2 Definition of “user communities”**

In this paper, “user community” is used to represent a logical grouping of users that can be seen as a unique interlocutor for all other actors. In our Grid context, a typical example of a user community is a Virtual Organization (VO), but this can also be extended to Virtual Research Communities (VRCs) or a specific scientific community federated around a given project.

### **1.3 Definition of “resource providers” and “service providers”**

Resource providers are the entities that provide user communities with access to computing and storage resources. They are grid resource centres or “sites” as described in EGI operational architecture definition [4]. Service providers are entities offering services that can be technical – e.g. core VO services, monitoring tools etc. – or not – e.g. support or expertise. France Grilles places itself both as a service provider and an operation centre as defined in [4].

### **1.4 Definition of “resource allocation”**

We consider “resource allocation” as a process involving different partners with the aim to provide resources and services to Grid users. The result is actual resources and services being provided, but also agreements being set up between providers and consumers. Involved actors are user communities, resource providers and hyper-structures such as Grid Infrastructures. “Allocation” is understood in this context, and should not be interpreted as “reservation”.

## **2 Identified needs and goals**

### **2.1 Improve service delivery to the French Community**

The French Grid Infrastructure “France Grilles”, like most of its counterparts, has been set up to answer specific scientific needs according to a national scientific policy. Overall supervision on resource allocation and distribution is highly desirable and should be done in respect to this. An allocation policy is

clearly needed to get a comprehensive and fair understanding of resource distribution. If there is a need to re-equilibrate this distribution of resources between different communities (VOs, VRCs, projects, etc.), this should be done according both to the needs and the overall scientific policy. This is an essential contribution to a better quality of service delivered to our users.

## **2.2 Measure what is done**

There has been no clear resource allocation policy so far within France Grilles: the current resource allocation model which was designed with the needs of EGEE projects should be revised to ensure the visibility and sustainability of the French Grid Infrastructure. Beyond that, there is a clear need of accountability. Especially, France Grilles needs to be able to:

- Assess how resources and services are delivered to the French community;
- Justify that resources delivered to international communities are not wasted, and that there is a return on investment.

## **3 Inspiration from existing resource allocation mechanisms**

### **3.1 High Performance Computing (HPC) world**

Resource allocation is an important aspect of all computing infrastructures, and High Performance Computing is no exception. In this particular domain, resource allocation is based on scientific evaluation, through a priori (e.g. evaluation of answers to a call for proposals) and a posteriori (regular review of supported demands) analyses. This is how resources are allocated in the GENCI [5] project, as explained as early as 2007 in [6] and reflected in the yearly activity report from 2009 onwards [7].

These concepts could be adapted to the Grid context. However this has to be done with care to take into account Grid specificities such as “free” access to resources and the absence of the concept of resource reservation.

### **3.2 Worldwide LHC Computing Grid (WLCG)**

The WLCG [8] resource allocation model is based on a principle of pledges: to answer LHC experiments’ needs, participating resource providers offer amounts of computing resources under the form of pledges under the supervision of a “Resource Scrutiny Group” on a yearly basis [9]. WLCG being the biggest user of France Grilles resources, it is utterly important to take this procedure into account.

### 3.3 Other national grids

Discussions and collaborations with the polish NGI PL-Grid [10] have led to sharing ideas and concepts about resource allocation at a national level. PL-Grid model is a resource allocation centric model [11] which makes extensive use of an SLA management tool, the Grid Bazaar [12]. Interactions with our Polish colleagues have already produced some of the ideas described in this paper. The use of a bazaar-like tool is equally one of the tracks we could follow in the future.

## 4 Definition of the strategy

### 4.1 Key principles

We propose the following principles as a basis for our resource allocation strategy:

- Decisions on how to allocate resources are made on both *a priori* and *a posteriori* analyses, the former allowing to agree on estimated needs and the latter focusing on measuring how much has been used
- New communities can join in and use resources without necessarily being filtered, provided their needs are reasonable (filtering is done above a given threshold in terms of how much the user asks – if asking for any precise amount of resources)
- Established user communities provide the scientific expertise needed to validate resource allocation above this threshold
- There is a unique point of contact for all users in demand of resources
- The complexity of the model is not visible to users
- The whole model allows to measure and report on resource usage for both new and established communities, either French or international

#### 4.1.1 Who are the user communities we have to consider?

There are various kinds of user communities using French NGI resources, spanning from international to regional, thematic or project-driven. Moreover, we are now considering Virtual Research Communities within EGI. These VRC will gather several VOs spread across several projects, countries and groups. Our needs and key principles lead to two clearly different use cases:

- Resource allocation to **new users** (not using the grid yet) and **French scientific communities**. Those communities might not be structured yet and can be identified by the project that federates them.
- Resource allocation to **established international communities**. These can be international VOs or VRCs

## 4.2 Allocating resources to new users and French communities

### 4.2.1 “A priori” analysis

In the overall scenario described on figure 1, a new user with a predefined project asks for resources, or simply expresses interest to join the Grid without any precise demand on the amount of needed resources. The request is handled by the NGI through a single point of contact that acts as a “broker”. The 3 basic questions to answer at this stage are:

1. Is there an existing VO on the Grid that could integrate this project/demand to its activities?
2. Is the user “grid aware”, e.g. is the project ready for grids, are all applications ported etc.?
3. Is the requested amount of resources above a given threshold, if any?

As shown on Fig.1, the result of the analysis can either be:

- **Rejection of the project** if it is considered non valid by the scientific committee
- **Redirection of the user:**
  - o **To the training activity** if it is felt the project has potential but is not grid-enabled or grid-focused;
  - o **To a better frame** (e.g. HPC) if it is felt the project is not a good use case for grids. A bi-directional process is the long-term aim here: e.g. HPC potential users could be redirected to grids if their need is better matched.
- **Project support through a VO based resource allocation agreement.** In this case, an existing VO accepts the new user as one of its members and applies its own policies with regards to how much resources this user can get from what is already available to the VO. Example: a new user with a project in biology will probably be redirected to biomed, who will then decide what place to give to this project within their activities.
- **Project support through an NGI based resource allocation agreement.** This is the case we present in details below.

The exact composition of the scientific committee deciding on demands above defined thresholds will soon be defined in the context of the first French User Forum in September 2011. This should certainly involve scientific coordinators from the user communities, under the NGI umbrella.

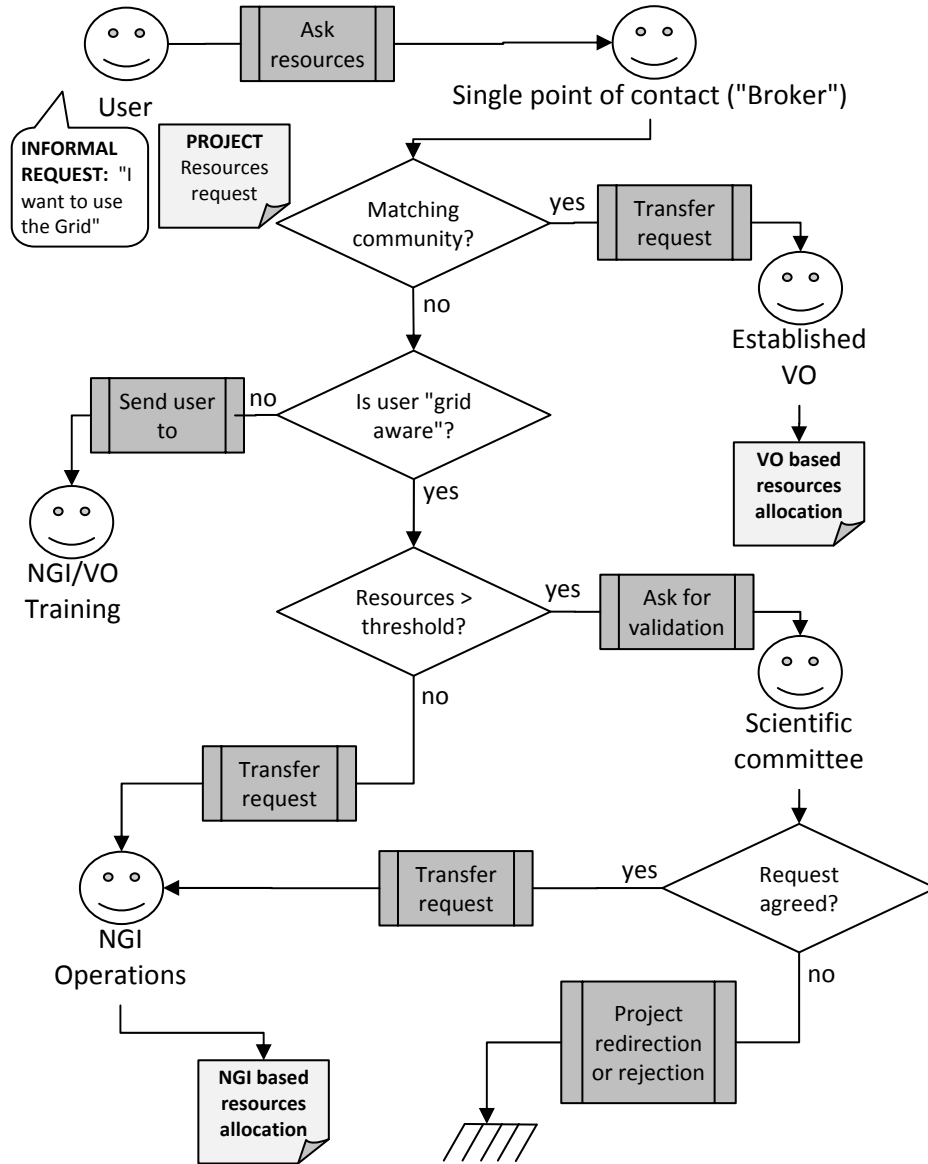


Fig. 1. A priori analysis for resource allocation requests from new users

The scientific committee also decides on the values for thresholds, as well as on any additional criterion needed for the evaluation of scientific validity of a given project and its interest for France Grilles.

#### 4.2.2 Project support through an NGI based resource allocation agreement

Depending on the scope of the resource allocation, each involved body should be able to decide at its own level. Agreements on “physical” resources (e.g. CPU) should be decided by sites, while agreements on services (e.g. support) should be made by the NGI. This is because the final decision should be taken by who controls the resources. Each site, as resource provider, has a different funding schema and is the best placed entity to commit to provide resources. At a higher level, the NGI doesn't have to control these resources but could just act as a relay.

The threshold principle applied within the a priori analysis can also be used to determine whether the project will be supported by the NGI through the creation of a new VO or through a catch-all VO.

A process proposal is described on figure 2.

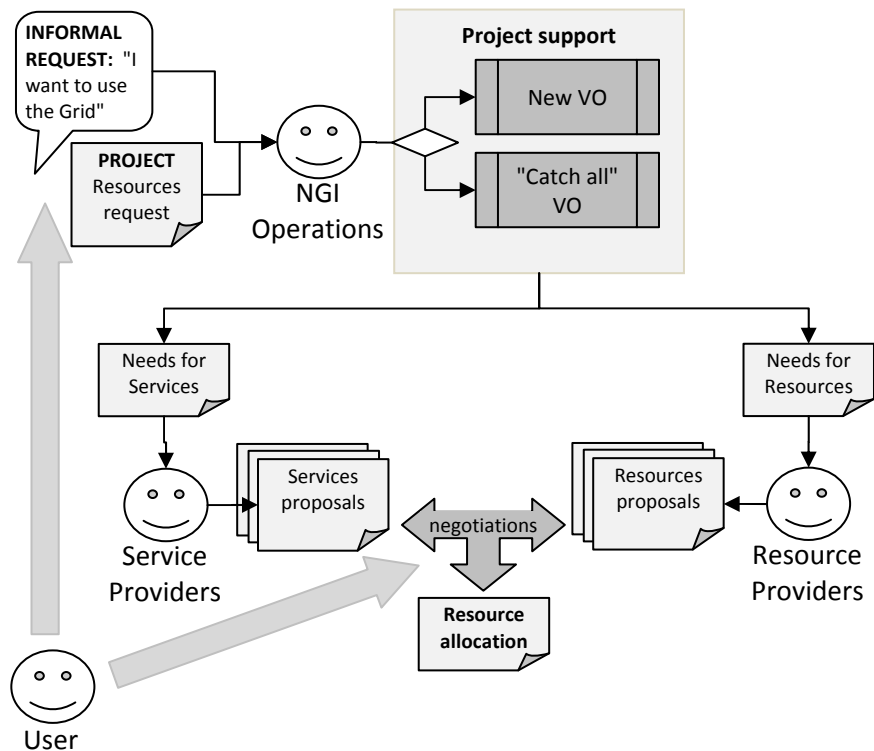


Fig. 2. Establishment of an NGI based resource allocation agreement



The result of the process is the establishment of a resource allocation agreement between the resource providers (sites), service providers (sites and/or NGI) and the user.

#### 4.2.3 “A posteriori” analysis

Resource usage verification for supported projects (i.e. those who have been allocated resources through an *a priori* analysis) leads to an *a posteriori* analysis of the initial application and possible review of new requests by projects.

The goal of this analysis is to:

- Assess the validity of the initial request
- Monitor the possible growth of the project, and take into account new resulting needs

At this stage, there is a need to define a second threshold in the amount of used resources above which the user/project which have been integrated to the catch-all community need to “emancipate” and start a new community.

In the case of a “VO based resource allocation” (see fig.1), this *a posteriori* analysis should allow to assess new needs for the considered VO. This will then be taken into account as part of the process of allocating resources to established communities.

Detailed workflows and implementation of an *a posteriori* analysis will be the subject of a deeper study in the months to come.

### 4.3 Allocating resources to established international communities

#### 4.3.1 Scope of the process

We aim that France Grilles resource allocation strategy include the case of international VOs whenever possible.

We are fully aware that some project driven communities (e.g. WLCG) already have a clear resource allocation mechanism: our goal is neither to temper with this nor to add an extra layer that would unnecessarily complicate the process. However, it is utterly important to provide a frame for international VOs who wish to negotiate resource allocation with NGIs and that our model is in line with WLCG’s.

#### 4.3.2 Proposed principle

We propose to deal with international VOs/VRCs in a similar way to French communities, by considering only the French part of this VO/VRC (e.g. LCG-

France for LHC VOs). From an NGI point of view, the interlocutor is then the representation of this VO/VRC in France.

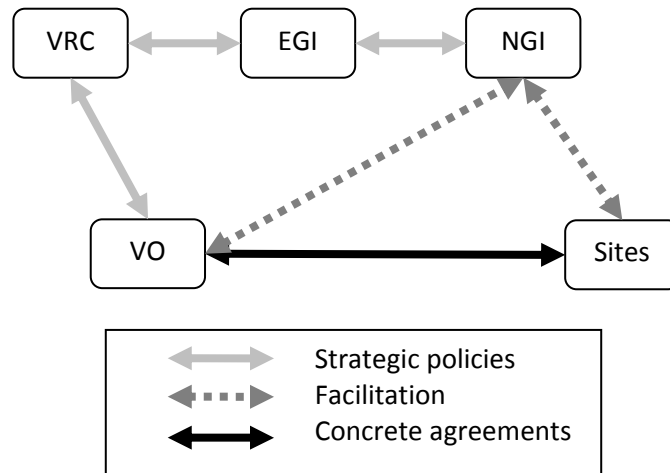


Fig. 3. Interactions between partners involved in resource allocation

From a VO point of view, France Grilles can act as a facilitator to reach agreements with sites. Depending on which granularity the VO considers convenient to deal with, agreements can be built at NGI level, or at site level. As an example, MoUs are already covered by the biomed VO for the latter category.

### 4.3.3 Measuring “French” usage of resources

One of the needs expressed at the beginning of this paper is to report on service delivery to the French community. If this is easy to do in the context of regional or national VOs, this is a more complex problem in the case of international ones.

Practical methods can be implemented to distinguish between French vs. foreign usage within a VO (e.g. calculating the ratio of certificate DNs issued by the French CA). Ideally, the request made to EGI to produce as a metric the percentage of usage par certificate DN per VO should help in that matter.

This approach can be limitative with regards to some usage of the grid such as pilot jobs in the scope of international communities; we expect though that this starting point in our estimation of the resource usage will improve with time.

## 5 Next steps

Jointly to this work on a national resource allocation strategy, we are currently in the implementation process of national VO for France Grilles. Establishing such a VO addresses the need for an easier integration of new users by establishing a VO supported nationwide and open to all. Our usage scenario is to add this national VO to the already available offer provided by local and regional VOs, so as to provide the French community with a larger spectrum of possibilities to answer their needs. This way we also build upon the existing structure and manpower set-up in regional grids to remain as close as possible from the end-user.

Resource allocation through the national VO can be seen as a possible implementation of the NGI based resource allocation agreement described earlier (see Fig.2).

As mentioned in section 4, a deeper study of the modalities of an *a posteriori* analysis is also needed to make any further progress. Part of our effort in the months to come will be dedicated to that.

Also, and as exposed earlier, the usage of a tool to monitor and follow negotiations between any resource providers and user communities is currently under study. This could lead to the set-up of an *a posteriori* usage dashboard and possibility to drive the process for *a priori* allocations for resource allocation through the national VO. Also such assessments could be used for the real time implementation of the resource allocation.

## References

1. EGI-Inspire web site, [www.egi.eu/projects/egi-inspire/](http://www.egi.eu/projects/egi-inspire/)
2. EGEE web site, [www.eu-egee.org](http://www.eu-egee.org)
3. France Grilles web site, [www.france-grilles.fr](http://www.france-grilles.fr)
4. Ferrari T.: EGI Operations Architecture, EU deliverable D4.1, <https://documents.egi.eu/public/ShowDocument?docid=218>
5. GENCI web site, [www.genci.fr](http://www.genci.fr)
6. Riviere C.: GENCI : Grand Equipement National de Calcul Intensif. Rencontre GENCI ORAP, 2007, Paris, <http://www.genci.fr/spip.php?article13>
7. Riviere C. : Rapport annuel 2009 de GENCI <http://www.genci.fr/spip.php?article92>
8. WLCG web site, [lcg.web.cern.ch](http://lcg.web.cern.ch)
9. WLCG MoU, Annex 9 "Rules of Procedure for the Resources Scrutiny Group (RSG)", <http://lcg.web.cern.ch/LCG/mou.htm>
10. PL-Grid web site, [www.plgrid.pl](http://www.plgrid.pl)
11. Szeplieniec T., Radecki M., Tomanek M.: A Resource Allocation-centric Grid Operation Model, in proceedings of the ISGC 2010 conference, Taipei, Taiwan (2010)
12. Bazaar Project Web Page, <http://grid.cyfronet.pl/bazaar>