France Grilles – Resource Allocation Strategy

Main authors: Gilles Mathieu, Hélène Cordier
Contributions from : Tristan Glatard, Emmanuel Merdernach, Monique Petitdidier, Frédérique Chollet, Yannick Legré
Last modified: 24/02/2011

Identified needs and goals

Improve service delivery to the French Community

Global control on resource allocation and distribution is highly desirable and should be done in respect to France Grilles' scientific policy. A 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 as well a prerequisite for delivering a high Quality of Service to our users.

Measure what is done

There has been no clear policy so far: the current resource allocation model which was designed with the needs of EGEE projects should be revised to ensure the visibility and sustainability of France Grilles. 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.

Definition of the strategy

Key principles

- 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)
- Established user communities provide the scientific expertise needed to validate resource allocation above this threshold
- France-Grilles defines 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

Who are the user communities we are talking about?

There are various kinds of 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, projects, countries and groups.

Our identified needs and our key principles then 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

Allocating resources to new users and French communities

"A priori" analysis

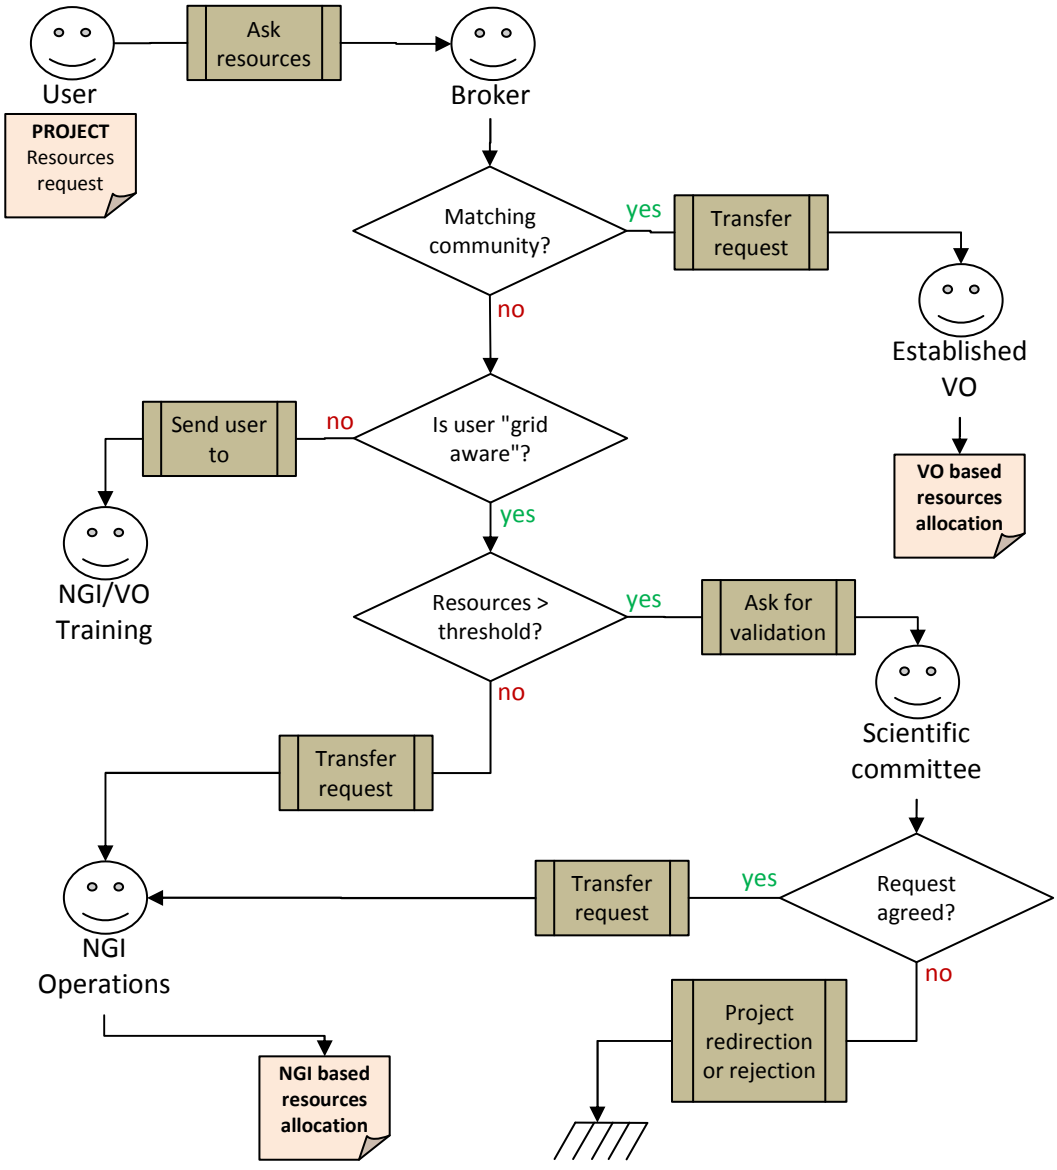

Fig.1 – A priori analysis for resource allocation requests from new users

In the overall scenario described on figure 1, a new user with a predefined project asks for 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 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?

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:**
  - **to the training activity** if it is felt the project has potential but is not grid-enabled or grid-focused
  - **to a better frame** (e.g. HPC) if it is felt the project is not a good use case for grids
- **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 remains to define, but this should certainly involve scientific coordinators from the user communities, under the NGI umbrella.
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.

Project support through an NGI based resource allocation agreement

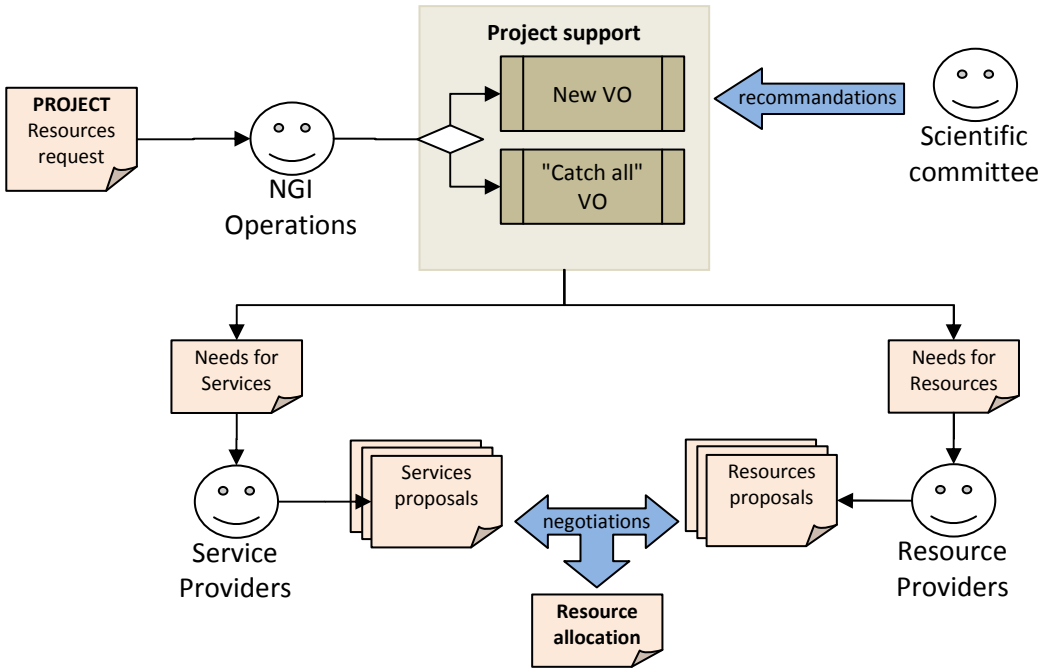

Fig.2 – Establishment of an NGI based resource allocation agreement

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.

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.

"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.

Allocating resources to established international communities

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.

Proposed principle

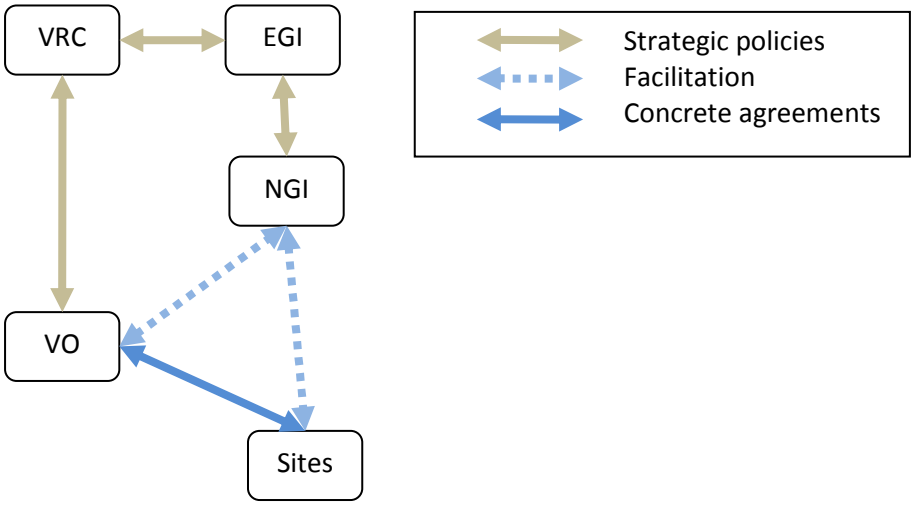

We propose to deal with international VOs/VRCs in a similar way to French communities, by considering only the French component of this VO/VRC. From an NGI point of view, the interlocutor is then the representation of this VO/VRC in France.
From a VO point of view, France Grilles can act as a facilitator 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.

Measuring "French" usage of resources

To distinguish between French vs foreign usage within a VO, a practical solution could be to multiply the total VO usage by the fraction of the VO DNs delivered by the French CA. For instance, if biomed consumed 100 hours and 27% of the DNs in biomed are French then we consider the French biomed usage to be 27 hours.

More discussion is needed to define a complete strategy in this context. To be continued at EGI User Forum in April.