

Report from the meeting of the IdG France Grilles International Advisory Committee

Date: 22nd March 2011

International Advisory Committee members present at the meeting:
Henri Bal, Ian Foster, Bob Jones (acting chairman), Domenico Laforenza, Richard Mount, Philippe Navaux, Ruth Pordes (by phone)
Apologies: Henri Casanova, Neil Geddes, Satoshi Matsuoka, Mitsuhisa Sato

General Comments

The reformulated International Advisory Committee held its first meeting in Lyon on 15-16 March 2011. The organization of this two day meeting had taken into account the recommendations from IdG Scientific Council in 2010 and provided for a more direct and in-depth communication between the IdG management and the committee members. The agenda and material presented during the meeting is here: <http://indico.in2p3.fr/conferenceDisplay.py?confId=5156>

Vision and Strategy

A vision has been presented which is compatible with the resources IdG has available. Elements of the strategic plan have been developed and these need to be brought together into a consistent proposition for all the stakeholders. Such elements would be yet more compelling if they were made more concrete in terms of the specific end state that they aspire to realize. The vision is, presumably, a world-class national infrastructure that can support the HTC needs of dozens of other science communities. Realizing this vision will then require tactical plans for: operating a large national infrastructure; expanding the user community; and expanding the resource base. Overall progress has been made with varying levels of success and impact in different domains. The work of IdG is dependent on relationships with a number of national bodies which are complicated and require management. A stated policy from the funding bodies on the use of distributed computing by the research community would enable IdG to operate more effectively in this national context.

A major activity of IdG during the last year has been the transition from EGEE to EGI at the European level. This transition has been well managed with no loss of quality of service or user engagement despite the reduced level of manpower available via European Commission (EC) funded projects. IdG developments are consistent with the EGI approach and France has provided valuable contributions to EGI. The production infrastructure has grown in proportions that are in line with the increased resources installed at the participating sites.

We believe that IdG could benefit from more careful analysis and measurement of the impacts of their work on different scientific communities. There are many data that could be cited: certainly publications, but also number of users, examples of scientific outputs, number of tasks executed, and students engaged. We heard various examples of campuses that had benefited from their local grid infrastructures (e.g. by engaging users who would not otherwise make use of HPC) and this information should be captured in some way.

The relationship with the Computer Science (CS) community, notably Grid5000 and ALADDIN, is positive and should be encouraged. The initiative of offering minor sums of funding to establish collaborations between the communities is seen as a success. Specific examples of transfer of technology between Grid5000 and production grids were highlighted (e.g. DIET, JSAGA etc.) as well as the CS publications coming from the Grid Observatory. This relationship is an advantage for France should be further developed.

Relations with supercomputer centres via GENCI are only just starting (notably via the earth and life science communities) with some joint project proposals submitted but not funded. We suggest that IdG builds on the CSCI recommendation for HPC and grids to collaborate to drive forward interaction. Such interaction could be implemented by focusing on a single application and supercomputer site which can be developed and act as a valuable use case for the wider research community.

The relationships with other projects (e.g. StratusLab) are seen as very beneficial as a means of exploring the benefits of new technologies, such as virtualization, and should be reinforced. This will require more work to analyse the results and develop a strategy for their introduction. We suggest that a small fraction of the equipment budget could be used to purchase access to commercial public cloud services for specific user applications on a trial basis. Given its relations with several projects and user communities, IdG has an opportunity to develop a leadership role for itself with respect to national policy for academic cloud computing.

We encourage IdG to produce and distribute written versions of the strategic plan and technical roadmap taking into account the recommendations of this report. The vision, strategy and roadmap should be clearly visible on the website and wiki pages.

User Support and Training

Excellent support has been provided to the High Energy Physics (HEP) community during this year and this should continue in the future since it represents the largest user community for IdG. Outreach to other communities has been an important activity and links to the LifeWatch, EPOS and ELIXIR ESFRI projects have been established. IdG should work with the national groups involved in these projects to facilitate their use of distributed computing services. Additional resources will be required in the user support domain if IdG is to be successful in converting such outreach initiatives into committed adopters of distributed computing services.

Some successes in the data management area have been achieved, notably with the NeuroLOG project, and we believe the tools and expertise in federating data sets are likely to prove popular with many user communities.

The usage data statistics are immature due to limited tooling and analysis. Such data is very important for reporting the impact of IdG to its stakeholders and merits more attention. Further investigations are necessary to be able to report in an authoritative way whether IdG is meeting its commitments. The use of graphical representations will certainly enhance the understandability of such metrics.

Training activities were presented but without any metrics. The website has information about training services and resources. The training of students is an important output which should be highlighted and included in the strategic plan.

Considerable effort has been put into international training but this should be complementary to domestic training efforts.

IdG is making some progress in engaging communities beyond high energy physics. To reach the point where a community can make optimum use of the national infrastructure, we estimate that 1 full-time person is needed to support each major user community. We recommend that IdG start with 2 user communities and ensure they receive sufficient support in order to be able to have a significant impact.

The current allocation of resources to training and user support needs to be analysed quantitatively. We suggest a holistic approach is adopted between the technical and user support to maximise the expertise and experience that exists in each area.

The plans for a catch-all national virtual organization appear useful for new, individual users. The services offered to users need to be described and marketed widely.

The users' opinions need to be systematically collected and analysed. The plans to organize a user forum in September will provide a good opportunity to gather such feedback and offer training to new users.

Summary of Recommendations

1. Integrate the vision and elements of the strategic plan into a consistent proposition for all stakeholders.
2. Produce and distribute written versions of the strategic plan and technical roadmap.
3. Promote the vision, strategy and plans for IdG via the website.
4. Produce easily understandable and consistent metrics to measure the performance and impact of IdG with respect to its objectives and commitments.
5. Further develop the fruitful relationship with the computer science community.
6. Focus the initial interactions with supercomputer centres by implementing one application in cooperation with a single supercomputer centre.
7. Consider using a small fraction of the equipment budget to purchase access to commercial cloud services for targeted user communities on a trial basis.
8. Identify two non-HEP communities to each be supported by a dedicated full-time member of the IdG team.
9. Quantitatively analyse the current allocation of resources for user support and training.
10. Collect users' opinions on a regular basis.

Finally we would like to thank all the staff of IdG and France Grilles for their hospitality and the hard work that went into the preparation for this review.