

Towards a European Infrastructure for e-Science

Enabling Grids for E-science (EGEE) consortium – August 2006

Introduction

Throughout Europe, Grid computing has been developing via a series of short-term projects with wide geographical coverage funded by national governments and the European Commission. Several communities are now successfully using these infrastructures as this technology opens new possibilities and scientific success often depends on them. To exploit fully the benefits that Grid computing can offer to European science and industry, it is necessary to put in place a more permanent structure to ensure the continued and sustainable availability of such computing facilities. This paper explores the motivation for such a structure, how it would benefit its stakeholders, what it could be expected to provide and where it would be positioned in relation to existing and future e-infrastructures.

Motivation

Grid computing will eventually make access to resources held on widely distributed computers as easy as access to resources on the user's own desktop. A Grid infrastructure that provides networking, computing and data resources in such a way that they are readily available to users regardless of their geographical location will improve the efficiency of scientific and industrial research and be of great benefit to the wider society. In addition, it will benefit those domains which traditionally were not users of large-scale computing facilities. Digital Libraries are one example of such a domain, although in this case the benefits could be mutual. The need for a sustainable e-infrastructure has already been recognised by the Organisation for Economic Co-operation and Development and the e-Infrastructure Reflection Group of the EU^{1,2}.

Many countries now have National Grid Initiatives (NGIs). NGIs, operating application-independent national e-infrastructures and supported by national funding and resources, help to define international standards and policies. In addition, through the 5th and 6th Framework Programmes, the European Commission has demonstrated its commitment to Grid computing as an emerging technology with the potential to contribute to many areas of the European economy, from improving scientific research to easing the digital divide. One of the first projects funded was the European DataGrid, completed in early 2004, which built a successful test-bed to demonstrate the scientific potential of Grids. This complemented the European DataTAG project, which established new records in long-distance data transfers via international networks. It was superseded by the Enabling Grids for E-science (EGEE) project, which has constructed production-quality infrastructure and built the largest multi-science Grid in the world, with over 200 sites.

Such efforts have firmly positioned Europe as a world leader in Grid computing. With related projects and national Grid programmes they have proved the success of Grid computing for the research community and as a valuable component of the future European Research Area (ERA). The technology is still developing, and EGEE-II features an expanded consortium with many more application domains and a host of related projects.

¹ Report on Grids and Basic Research Programmes. OECD Global Science Forum, Sydney, Australia, September 2005. <http://www.oecd.org/dataoecd/30/36/36213997.pdf>

² Malcolm Read, A European vision for a Universal e-Infrastructure for Research. e-IRG meeting, London, UK, December 2005. http://www.e-irg.org/meetings/2005-UK/A_European_vision_for_a_Universal_e-Infrastructure_for_Research.pdf

Despite these successes, the current state of European Grid computing is not sustainable. There is much enthusiasm among user communities, but efforts such as EGEE which are funded on two-year funding cycles are unable to provide sufficient guarantees for the long-term availability of services necessary for mission-critical applications and tasks. The size of the EGEE-II consortium is also approaching the limit for effective management and coordination of a single project, but such large-scale coordination is necessary in order to provide production-quality infrastructure. This situation would be greatly aided if representation at a future body were at the level of NGIs. Short-term project-based funding cycles also risk dispersing the skill base that the FP5 and FP6 Grid projects have built up, a skill base that must have long-term sustainability if the potential of Grid computing is to be fully realised in the years to come.

Stakeholders

In considering new models for the future of the European e-Science infrastructure, the requirements of a large number of stakeholders must be taken into consideration. These include:

- **User communities** – From the users' point of view the complexity of the Grid must at all times be minimised, although this is not easy when accounts must be held with a number of services and entities that may not communicate with one another. The situation would be aided by the centralisation of infrastructure oversight, hence promoting the unification of Grid services irrespective of the underlying organisational structures and systems. Further developments are required to ease the access to Grids for users and small collaborations which do not have sufficient local expertise to use complex systems ('computing out of the wall'). Another service beneficial to users would be a centrally-administered set of computing resources that could enable new users to gain access to Grid technology.
- **Existing Grid projects** – Within Europe the largest Grid projects, such as EGEE, already act as hosts and incubators for a range of smaller projects but, as the Grid community grows, such activities will increasingly stretch their resources. These projects would significantly benefit from the long-term existence of some central body able to provide them with support and to promote innovation and minimise duplication of effort.
- **National Grid Initiatives** – The development of Grid computing has in many countries reached the point where there is a single concerted effort overseeing the national Grid infrastructure and strategy. These National Grid Initiatives (NGIs) provide a single point of contact, reducing the management and organisational overhead of international cooperation, and their existence has been identified as a key step in the development of a European Grid Infrastructure.
- **Industry** – Industrial interest in Grid computing is inevitably based around potential future markets and profits. At present, industry is involved in a wide range of projects for enterprise Grids and is starting to see the need and benefits for shared Grid infrastructures between many business partners. Industry is also involved with efforts such as NESSI and the Global Grid Forum (following a recent merger with the Enterprise Grid Alliance), but in Europe they lack a single point of contact.
- **Resource providers** – Currently resource providers, be they NGIs or individual organisations, must be heavily involved in the management of the infrastructure and considerable work is required to maintain a Grid site. This situation would be greatly improved by the development of agreed Europe-wide practices for resources management. It would also benefit from a centrally certified and tested middleware distribution system that would enable sites to "mix and match" components of assured quality which best fit their requirements. Centralised infrastructure coordination would

also reduce overheads and encourage new centres to contribute resources, either directly, or via their NGIs where these exist.

- **Policy and standards bodies** – Research is increasingly multi-national and the investment that has already been made will only have its maximum pay-off when Grids can interoperate, regardless of their country of origin or scientific discipline. At present there are a number of policy bodies concerned with Grid technology in Europe, each of which must deal with many individual projects, be they thematic, national or trans-European. The existence of a single body with appropriate consultative mechanisms to represent the infrastructure providers, application groups and European Grid community in general would accelerate the development and adoption of internationally agreed standards and acceptable use policies that are crucial for long-term development and the wide adoption of Grid computing.
- **The public** – While Grid computing is not yet developed to the point where public involvement is possible, the ongoing move to a service-orientated IT sector suggests that such distributed computer systems soon will be of public interest. In order to reduce barriers to future public participation it will be necessary to provide the kind of seamless service that the public has come to expect through use of broadband networking and the growing number of on-demand services.
- **National/EU/funding agencies** – The EU and a range of private and national funding bodies have invested heavily in Grid technology over the past decade. At present their investment is at risk, in part because of the funding models they necessarily use. This danger occurs as Grid technology matures and disparate efforts coalesce into a smaller number of key initiatives which require longer-term stability. The formation of a structure to manage Grid computing in Europe would secure existing investments and maximise their return over the coming years.

Proposal and Objectives

To deal with the issues raised here in such a way as to fulfil the needs of the stakeholders, a new paradigm for the organisation, coordination and management of European Grid computing must be considered. Thus, a body should be set up to handle these issues Europe-wide and must be funded on a long-enough time-scale. Parallels can be drawn with the condition of scientific research networks in the early 1990s, when the creation of DANTE and the TEN34/155 and later GÉANT projects united disparate efforts and guaranteed the long term availability of high speed networking for the ERA. A comparable change in the field of Grid computing would help to capitalise on the European Union's considerable investment in this area and provide Europe with a multilayered approach to scientific computing, with the GÉANT network supporting both Grid and supercomputing infrastructures (such as that provided by the DEISA project).

Achieving this will require a new organisational structure with two new kinds of entities. The first involves the formation of National Grid Initiatives (NGIs) which unite efforts within each country, providing a single point of contact for coordinated efforts and reducing the management and organisational overhead of international cooperation around Grid computing. While some countries or regions already have analogous organisations, the existence of a full European network of NGIs has been identified as a key step in the development of the European Grid Infrastructure. These national bodies would be complemented by a second structure to coordinate issues on a European scale, including operations management, policy, standards and middleware testing. Such a structure addresses the key issue of long-term sustainability of the European Grid Infrastructure, ensuring the availability of Grid services to researchers from all backgrounds and disciplines as we move further into the 21st century.

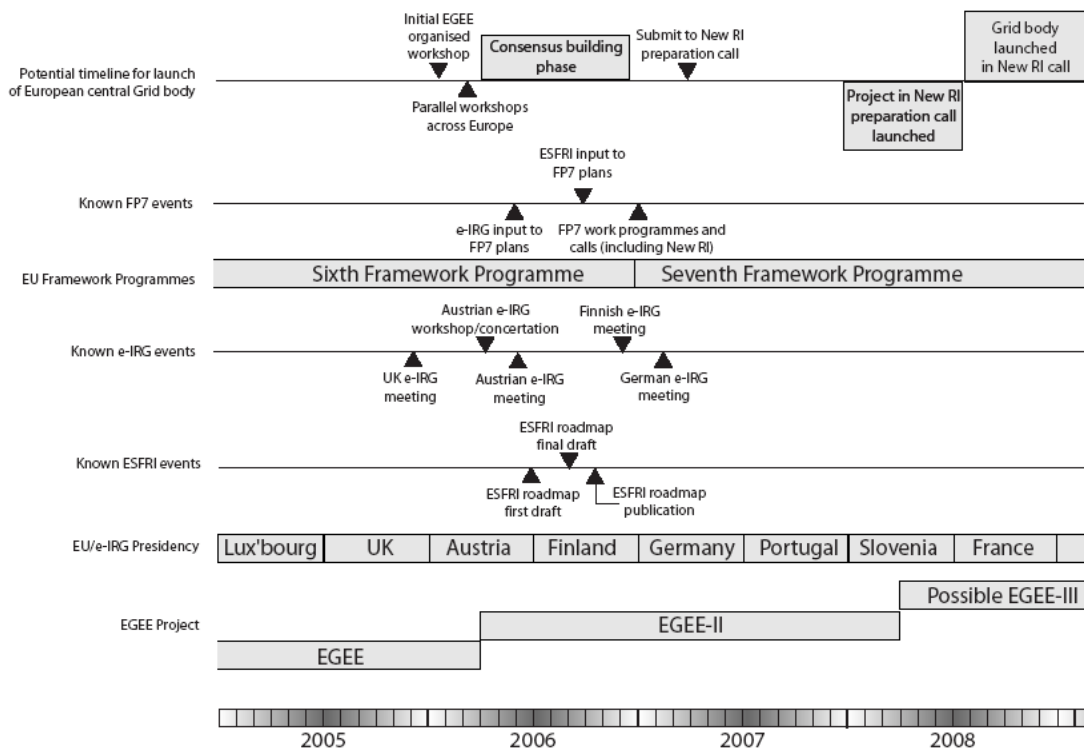
From consideration of the various stakeholders involved, a number of core objectives for any future centralised European Grid Computing hub become clear:

- The key functionality to be provided by a potential Grid organisation is the coordination of the operation and oversight of the basic European Grid infrastructure, a task presently carried out separately by each participant. Centralisation of these services, in close cooperation with National and Regional Bodies, would help bind the different European Grids into a seamless whole, in line with the long-term vision and promoting a truly virtualised service from the point of view of the user. This would also ease the interoperability of Grid computing with complementary efforts in High Performance Computing such as the DEISA project.
- Complementing the operations coordination, such a central hub would also provide a testing, certification, and validation service for the infrastructure. This task would help build the brand of the proposed body into the European quality mark for Grid computing. In combination with operations, it would also help ensure Europe-wide Quality of Service, a problematic issue in such a heterogeneous computing environment. This functionality is in line with the 'service' model for European Grid computing that runs through many of the concepts presented here, and in a similar vein a centralised structure would support a range of the activities commonly carried out in European Grid projects, such as training and application porting. It is unlikely that any central body will be able to provide such services directly to each individual project due to their diversity, and equally the staffing levels required would be prohibitive. However, what it can do is to support efforts in these areas with best practice guidelines and foundation material, and by providing a forum for exchange of information between projects.
- As Grid technology gathers momentum there will be additional scientific domains wishing to exploit its power as a research tool. Many of these newcomers will not previously have acquired many computing resources to contribute to a Grid, but nonetheless wish to explore its potential benefits. Thus, the possibility should be considered of establishing a small set of resource centres managed by the NGIs and centrally coordinated to provide a backbone of computing resources distributed throughout the ERA. This would be of particular benefit to introduce such research communities to the possibilities offered by Grid computing without such a facility having to be provided separately by each NGI.
- The unifying effect of a central hub for European Grid Infrastructure would have an impact not only within the community, improving Quality of Service and reducing the barriers to entry, but would also help to rationalise relations with those outside of the European Grid community. Such a hub could provide a concerted voice on standards bodies that would help accelerate their development, and hence the overall uptake of Grid computing. It would also simplify relations with communities in related technologies, such as networking and High Performance Computing, again maximising synergy and minimising duplication of work. This highly visible activity would address a number of common criticisms of Grid computing from the industrial sector, which often is disturbed by the lack of consensus and standards despite considerable work in this area.
- Finally, a central hub could provide strong communication links, operating internally and externally, both building community spirit among Grid users and publicising European Grid efforts to a wider audience. Internally this will aid in building synergies and improving productivity, while externally it will act as a multiplier for the impact of the overall organisation, by maximising its visibility on the world stage. Such efforts in the United States already benefit from more centralised communication due to the smaller number of projects (despite their large consortia) and the small number of funding sources. Similar benefits should be achievable in Europe.

Timeline/Roadmap

For any new Grid structure to be formed the schedules of existing related groups must be considered. The diagram summarises the projected timelines of the EGEE project, the e-Infrastructures Reflection Group (e-IRG), the European Strategy Forum on Research Infrastructures (ESFRI) and EC framework programmes over the period 2005-2008. A potential timeline for the new Grid structure proposed here is also included.

The e-IRG hosts a Sustainable e-Infrastructures Task Force which has written a paper including recommendations on the future of e-Infrastructures and submitted to the EC at the end of April 2006. ESFRI will draw up a roadmap for the future development of research infrastructures in Europe, due to appear in Autumn 2006. Both of these outputs will feed into the EC's decisions about funding in FP7. At present the plan is that FP7 will include a call specifically for 'New Research Infrastructures' (such as a sustainable Grid infrastructure) and this will include an initial funded phase in which such large undertakings can be planned before a final proposal is submitted. During this initial phase it is proposed that current funding streams continue to allow existing shorter-term projects, such as the development aspects of EGEE-II, to complete their research and development work.



External Relations

The formation of a central hub for the coordination of the European Grid Infrastructure would have a rationalising and structuring effect on the relationship between the European Grid community and the outside world. Relevant external groups include:

- **Non-European Grid projects** – The European Grid Infrastructure would be open to new resources and applications from the rest of the world; European Grids already

have a history of collaboration with groups such as OSG³ in the United States and NAREGI⁴ in Japan. Having a single voice to speak to these projects at a high level would improve communication and avoid such groups having to relate to the large number of individual European Grid efforts. Such collaboration would contribute to the eventual formation of global Grid consortia, much as in the field of Policy Management Authorities (PMAs), where regional PMAs have formed the International Grid Trust Federation⁵ to coordinate their efforts.

- **IT and standards bodies** – This would include groups such as the Open Grid Forum, the Institute of Electrical and Electronics Engineers, the International Telecommunication Union and the International Organization for Standards. A central body acting as a conduit for information would be beneficial for the creation and adoption of Grid standards.
- **Governments and funding agencies** – Having a single body capable of speaking for the broader community would help to maintain the high profile of Grid technology within the EU, both centrally and in the member states. In cooperation with the NGIs, it would also be able to lobby for the continued investment in the technology in a way that would benefit all stakeholders.
- **Industry** – The European and global IT industries are at present interested in Grid computing, as evidenced by their participation in a number of initiatives. A central body could provide appropriate information and direct interested parties to initiatives in their technical field or geographical region.
- **The Public and the media** – A central body would pave the way for public Grid use in the future by establishing a trusted brand for Grid computing in Europe. It would act in much the same way suggested for industry, acting as a meeting place for members of the Grid community, the media and the public. Such activities are especially appropriate at the present time, when Grid technology has a rapidly increasing user community and its value can clearly be demonstrated.

Implementation Issues

The ideas and objectives for a European Grid body presented here have been introduced from the point of view of stakeholders. However, setting up an organisation to fulfil these objectives will entail a number of practical issues, including those listed here:

- The stakeholders would consist of the NGIs and other resource providers, user communities and the European Commission. Initially not all potential stakeholders may be in a position to join, but the structure should be open to them becoming affiliated later.
- The proposed body will need to be set up as a legal entity and a possible model to follow would be that of DANTE with a consortium involving all NGIs.
- A governing board would be necessary to provide oversight and all NGIs should be represented on this. The EU involvement would be via annual reviews with EC appointed independent reviewers.
- A possible financial model would be for the stakeholders (mainly the NGIs and EU) to fund the central organisation.

³ Open Science Grid, www.opensciencegrid.org

⁴ National Research Grid Initiative, Japan, www.naregi.org

⁵ International Grid Trust Federation, www.gridpma.org

The above issues would be tackled during a short-term preparatory phase involving all the stakeholders and funded by the E.C. The outcome will depend on the makeup of the consortium, their needs and requirements.

In looking at organisational and funding models applicable to this situation the scheme used in the networking community for the GÉANT project clearly seems the most appropriate. Making such an analogy would liken the central hub to the DANTE organisation and the National Grid Initiatives to NRENs (National Research and Education Networks). Such a model certainly fulfils some of the requirements for a European Grid structure, but differences between the two communities (for instance the difference in the basic unit of operation, a packet or Grid job) must be fully understood if something close to this model is to be adopted.

Finally, CERN (The European Organization for Nuclear Research) has led several Grid projects and now leads EGEE. The lab's interest in a European Grid Organisation has been expressed in a position paper in June 2005, which started the discussion at a European level. The pioneering work performed by the High Energy Physics community has been extended to science in general and CERN believes that an independent organisation is required to assure the long-term future of Grids in Europe. Given CERN's unique experience in international collaboration, CERN offers to help setting up the new organisation and to host it initially, but believes strongly that CERN should become one of many stakeholders and clients at the same level as the NGIs in the future.

Conclusions

A new model for the coordination and operational oversight of the emerging European Grid infrastructure has been presented. This new model is necessary to serve the increasing number of stakeholders and to secure the technology's long-term sustainability and growth. The consortium presenting this paper is that of the EGEE project, as this represents the largest Grid infrastructure and a significant proportion of the Grid expertise in Europe. However, this is not a closed group and a sustainable European Grid infrastructure with central coordination will require broad consensus among all stakeholders. This paper sets out a number of key issues for discussion as such a broad-based consortium is formed.

Already a few of these key issues stand out and should be resolved at an early stage to promote the success of the overall venture. For example, there is a need to understand and synchronise with national and EC funding models and to investigate whether the EC is able to fund a sustainable Grid infrastructure in the same way as it does for networking. Following on from this, for maximum benefit to each region and to make administration and management of an integrated infrastructure possible, National Grid Initiatives or some equivalent bodies must be formed in all parts of Europe. While such bodies mostly exist in regions which are already well advanced in Grid adoption and development, care must be taken to ensure that all European nations are able to participate.

By acting on the issues and needs identified here, and building consensus on those less-well defined, Europe will be able to secure its position as the leading region for Grid computing for the years to come.