



Executive Summary

The executive summary describes the motivation for the EGEE (Enabling Grids for eScience in Europe) proposal recently made to the EU and the background activity behind this initiative

Background The state of computer and networking technology today makes the seamless sharing of computing resources on an international or even global scale conceivable. For scientific applications, the vision is that such computing Grids will integrate large, geographically distributed computer clusters and data storage facilities, and provide simple, reliable and round-the-clock access to these resources. The benefits that will flow include large increases in both the peak capacity and the total computing power delivered to various scientific projects, as well as new ways for scientific communities to share and analyse very large data sets. These benefits will translate into an increase of both the quality and quantity of scientific output in a broad spectrum of compute-intensive fields ranging from


bioinformatics and climate simulation to the nanoscale design of new materials and integration of large engineering projects involving many partners. Based on pioneering work in the US and in Europe, software toolkits for distributed computing – such as Globus, Condor, and Unicore – are available. As a result, a number of projects have demonstrated first early results for various aspects of computer Grids. Europe has achieved a prominent position in this field, in particular for its success in establishing a functional Grid testbed comprising more than 20 centres in Europe, in the context of the European DataGrid (EDG) project. Individual countries, such as the UK, France, and Italy, have developed comprehensive “e-Science” programs that

rely on emerging national computing Grids to deliver unprecedented computing resources to science. However, as yet, there are no real production-quality Grids that can offer continuous, reliable Grid services to a range of scientific communities.

The EGEE Vision EGEE (Enabling Grids for E-Science in Europe) aims to integrate current national, regional and thematic Grid efforts, in order to create a seamless European Grid infrastructure for the support of the European Research Area. This infrastructure will be built on the EU Research Network GEANT and exploit Grid expertise that has been generated by projects such as the EU DataGrid project, other EU supported Grid projects and the national Grid initiatives such as UK e-Science, INFN Grid, Nordugrid and the US Trillium (cluster of projects).

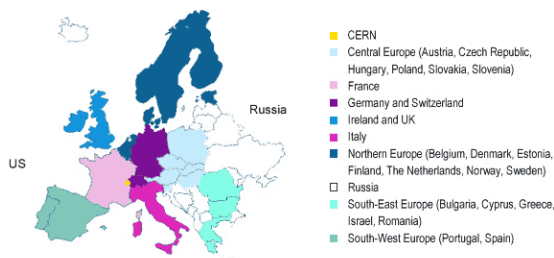
The EGEE vision is that this Grid infrastructure will provide European researchers in academia and industry with a common market of computing resources, enabling round-the-clock access to major computing resources, independent of geographic location. This infrastructure will support distributed research communities, including relevant Networks of Excellence, which share common Grid computing needs and are prepared to integrate their own distributed computing infrastructures and agree common access policies. The resulting infrastructure will surpass the capabilities of local clusters and individual supercomputing centres in many respects, providing a unique tool for collaborative compute-intensive science (“e-Science”) in the European Research Area. Finally, the infrastructure will provide interoperability with other Grids around the globe, including the US NSF Cyberinfrastructure,

contributing to efforts to establish a worldwide Grid infrastructure. The scope of the project is illustrated in the figure.



Schema of the evolution of the European Grid infrastructure from two pilot applications in high energy physics and biomedical Grids, to an infrastructure serving multiple scientific and technological communities, with enormous computer resources. The applications and resource figures are purely illustrative. The EGEE project covers Year 1 and 2 of a planned four year programme.

EGEE has been proposed by experts in Grid technologies representing the leading Grid activities in Europe. The process of developing this project has led to a structuring of the European Grid community into ten partner regions or “federations”. A significant structuring effect due to EGEE is already apparent, as several of these partners have begun integrating regional Grid efforts in order to provide coordinated resources to the EGEE project. In addition, US representatives are participating as EU unfunded partners in the project, and are considering establishing a US EGEE federation. Participation of Japan and the Asia-Pacific region is considered desirable and will be pursued.



The EGEE federations

EGEE is a two-year project conceived as part of a four-year programme. Major implementation milestones after two years will provide the basis for assessing subsequent objectives and funding needs.

Given the service oriented nature of this project, two pilot applications areas have been selected to guide the implementation and certify the performance and functionality of the evolving European Grid infrastructure. One is the Large Hadron Collider Computing Grid (LCG: www.cern.ch/lcg), which relies on a Grid infrastructure in order to store and analyse petabytes of real and simulated data from high-energy physics experiments at CERN. The other is Biomedical Grids, where several communities are facing equally daunting challenges to cope with the flood of bioinformatics and healthcare data.


Given the rapidly growing scientific needs for a Grid infrastructure, it is deemed essential for the EGEE project to “hit the ground running”, by deploying basic services, and initiating joint research and networking activities before the formal start of the project. The LCG project will provide basic resources and infrastructure already during 2003, and Biomedical Grid applications will be planned at this stage. The available resources and user groups will then rapidly expand during the course of the project. To ensure that the project ramps up rapidly, project partners have agreed to begin providing their unfunded contribution prior to the official start of the project.

The EGEE Mission In order to achieve the vision outlined above, EGEE has a three-fold mission:

1. To deliver production level Grid services, the essential elements of which are manageability, robustness, resilience to failure, and a consistent security model, as well as the scalability needed to rapidly absorb new resources as these become available, while ensuring the long-term viability of the infrastructure.
2. To carry out a professional Grid middleware re-engineering activity in support of the production services. This will support and continuously upgrade a suite of software tools capable of providing production level Grid services to a base of users which is anticipated to rapidly grow and diversify.
3. To ensure an outreach and training effort which can proactively market Grid services to new research communities in academia and industry, capture new e-Science

requirements for the middleware and service activities, and provide the necessary education to enable new users to benefit from the Grid infrastructure.

Reflecting this three-fold mission, EGEE is structured in three main areas of activity: **services, middleware re-engineering and networking**. Key aspects for each of these areas are summarised in separate sections below. It is essential to the success of EGEE that the three areas of activity should form a tightly integrated “Virtuous Cycle”, illustrated in the figure. In this way, the project as a whole can ensure rapid yet well-managed growth of the computing resources available to the Grid infrastructure as well as the number of scientific communities that use it. As a rule, new communities will contribute new resources to the Grid infrastructure. This feedback loop is supplemented by an underlying cyclical review process covering overall strategy, middleware architecture, quality assurance and security status, and ensuring a careful filtering of requirements, a coordinated prioritization of efforts and maintenance of production-quality standards.



The “Virtuous Cycle” for EGEE development

The Stakeholder **Perspective** In order to convey the scope and ambition of the project, this section presents the expected benefits for EGEE stakeholders and an outline of the procedure for the stakeholder to participate in EGEE. The key types of EGEE stakeholders are *users, resource providers, and industrial partners*.

a) EGEE Users

Once the EGEE infrastructure is fully operational, users will perceive it as one unified large scale computational resource. From the user perspective, the complexity of the service organisation and the underlying computational fabric will remain invisible. The benefits of EGEE from the user perspective include:

- **Simplified access** – Today most users have accounts on numerous computer

systems at several computer centres. The resource allocation procedures vary between the centres and are in most cases based on applications submitted to each centre or application area management. The overhead involved for a user in managing the different accounts and application procedures is significant. EGEE will reduce this overhead by providing means for users to join virtual organisations with access to a Grid containing all the operational resources needed.

- **On demand computing** – By allocating resources efficiently, the Grid promises greatly reduced waiting times for access to resources.

- **Pervasive access** – The infrastructure will be accessible from any geographic location with good network connectivity, thus providing regions with limited computer resources access on an as-need basis to large resources.

- **Large scale resources** – Through coordination of resources and user groups EGEE will be able to provide application areas with access to resources of a scale that no single computer centre can provide. This will enable European researchers to address previously intractable problems in strategic application areas.

- **Sharing of software and data** – By providing a unified computational fabric the EGEE will allow wide spread user communities to share software and databases in a transparent way. The EGEE will act as the enabling tool for European collaborations, building and supporting new virtual application organisations.

- **Improved support** – By making use of the expertise of all the partners EGEE will be able to provide a support infrastructure that includes in depth support for all key applications and around the clock technical systems support for GRID services.

A potential user community will typically come into contact with EGEE through one of the many outreach events supported by the Dissemination and Outreach activity, and will be able to express their specific user requirements via the Applications Identification and Support Activity. After negotiating access terms, which will depend, amongst other things, on the resources the community can contribute to the Grid infrastructure, users in the community will

receive training from the User Training and Induction activity. From the user perspective, the success of the EGEE infrastructure will be measured in the scientific output that is generated by the user communities it is supporting.

b) Resource Providers

EGEE resources will include national GRID initiatives, computer centres supporting one specific application area, or general computer centres supporting all fields of science in a region. The motivation for providing resources to the EGEE infrastructure will reflect the funding situation for each resource provider. EGEE will develop policies that are tailored to the needs of different kinds of partners. The most important benefits for resource providers are:

- **Large scale operations** – Through EGEE a coordinated large scale operational system is created. This will lead to significant cost savings and at the same time

improved level of service provided at each participating resource partner. Through EGEE, the critical mass needed for many support actions can be reached by all participating partners.

- **Specialist competence** – By distributing service tasks among the partners EGEE will make use of the leading specialists in Europe to build and support the infrastructure. The aggregate level of competence obtained is a guarantee for the success of the EGEE project. In this sense the Grid is used to connect distributed competence just as much as it is connecting distributed computational resources.

Each participating centre and its users will thus have access to experts in a wide variety of application and support fields.

- **User contacts** – The EGEE distributed support model will allow for regional adaptation and close contacts with regional user communities. The existence of regional support is of fundamental importance when introducing new users and user communities with limited previous experience of computational techniques. A resource partner in EGEE will become much more attractive as a collaboration partner on the regional level by representing the large scale EGEE infrastructure.

- **Collaborations among resource partners** – It is foreseen that several partners

within the EGEE framework will form collaborations and launch development and support actions not included in the present proposal. This will lead to cost sharing of R&D efforts among partners and in the longer perspective allow for specialization and profiling of participating partners to form globally leading centres of excellence within EGEE. These benefits motivate the many partners that support the EGEE proposal already, representing aggregate resources of over 17000 cluster nodes.

EGEE builds on the integration of existing infrastructures in the participating countries, in the form of national GRID initiatives, computer centres supporting one specific application area, or general computer centres supporting all fields of science in a region. The motivation for providing resources to the EGEE infrastructure depends on the mission and funding situation for each of the resource partners. A new resource provider will typically approach EGEE through contact with the Regional Operations Centres. Specific policy and contractual issues for a given resource provider will be dealt with by dedicated staff in the Operations Management Centre, based on general guidelines defined and regularly reviewed by the Project Executive Board, with advice from the Project Management Board, and reviewed regularly.

c) Industrial Partners

The driving force for EGEE is scientific applications, and the current partners represent publicly funded research institutions and computer resource providers from across Europe. Nevertheless, it is envisaged that industry will benefit from EGEE in several ways:

- **Industry as partner** – Through collaboration with individual EGEE partners, industry has the opportunity to participate in specific activities where relevant skills and manpower are available, thereby increasing know-how on Grid technologies.

- **Industry as user** – As part of the networking activities, specific industrial sectors will be targeted as potential users of the installed Grid infrastructure, for R&D applications. The pervasive nature of the

Grid infrastructure should be particularly attractive to high-tech SMEs, because it brings major computing resources – once only accessible to large corporations – within grasp.

• **Industry as provider** – Building a production quality Grid will require industry involvement for long-term maintenance of established Grid services, such as call centres, support centres and computing resource provider centres. The EGEE vision also has inspiring long-term implications for the IT industry. By pioneering the sort of comprehensive production Grid services which are envisioned by experts – but which at present are beyond the scope of national Grid initiatives – EGEE will have to develop solutions to issues such as scalability and security that go substantially beyond current Grid R&D projects. This process will lead to the spin off of innovative IT technologies, which will have benefits for industry, commerce and society going well beyond scientific computing. Major initiatives launched by several IT industry leaders in the area of Grids and Utility computing emphasize the economic potential of this emerging field.

Industry will typically come in contact with EGEE via the Industry Forum organised by the Application Identification and Support activity, as well as more general dissemination events run by the Dissemination and Outreach activity. Interested companies will be able to consult about potential participation in the project with the Project Director and with regional representatives on the EGEE Project Management Board. As the scope of Grid services expands during the second two years of the programme, it is envisaged that established core services will be taken over by industrial providers with proven service capacity. This service would be provided on commercial terms, and selected by a competitive tender. Service Activities The Service Activities will create, operate, support and manage a production quality European Grid infrastructure which will make resources at many resource centres across Europe accessible to user communities and virtual organisations in a consistent way according to agreed access management policies and service level agreements, while maintaining an overall

secure environment. These activities will build on current national and regional initiatives such as the UK e- Science Grid, the Italian Grid, and NorduGrid, as well as infrastructures being established by specific user communities, such as LCG.

The structure of the Grid services will comprise: EGEE Operations Management at CERN; EGEE Core Infrastructure Centres in the UK, France, Italy and at CERN, responsible for managing the overall Grid infrastructure; Regional Operations Centres, responsible for coordinating regional resources, regional deployment and support of services. The basic services that will be offered are: middleware deployment and installation; a software and documentation repository; Grid monitoring and problem tracking; Bug reporting and knowledge database; Virtual Organization (VO) Services; Grid Management Services. Continuous, stable Grid operation represents the most ambitious objective of EGEE, and requires the largest effort.

Middleware re-engineering activities

The current state-of-the-art in Grid Computing is dominated by research Grid projects that aim to deliver test Grid infrastructures providing proofs of concept and opening opportunities for new ideas, developments and further research. Only recently there has been an effort to agree on a unified Open Grid Services Architecture (OGSA) and an initial set of specifications constituting the Open Grid Service Infrastructure that set some of the standards in defining and accessing Grid services. Building a European Grid infrastructure based on robust components is thus becoming feasible. However, this will still take a considerable integration effort in terms of making the existing components adhere to the new standards, adapting them to evolution in these standards, and deploying them in a production Grid environment. The middleware activities in EGEE focus primarily on re-engineering existing middleware functionality, leveraging the considerable experience of the partners with the current generation of middleware. Based on experience, geographic co-location of development staff is essential, and therefore these activities are based on tightly-knit teams concentrated in a few

major centres with proven track records and expertise.

Networking activities The networking activities in EGEE aim to facilitate the induction of new users, new scientific communities and new virtual organisations into EGEE community. EGEE will develop and disseminate appropriate information to these groups proactively, and take into account their emerging Grid infrastructure needs. The goal is to ensure that all users of the EGEE infrastructure are well supported and to provide input to the requirements and planning activities of the project. Specific activities included in the EGEE proposal are: Dissemination and Outreach; User Training and Induction; Application Identification and Support; Policy and International Cooperation. The Application Identification and Support Activity has three components, two Pilot Application Interfaces – for high energy physics and biomedical Grids – and one more generic component dealing with the longer term recruitment of other communities.

Conclusions The EGEE project has successfully concluded the negotiation of the FP6 Research Infrastructure contract with the EU at the end of October and expects to start operations early Spring 2004. More than 160 researchers are gathering to participate in the various activities of the project and several new job positions have been opened by most of the seventy EGEE partners. If the first phase of the project will deliver by 2006, according to the expectations, a production quality Grid infrastructure for the European Research Area and the international scientific community more in general, it is planned to propose a second phase to extend both the geographical coverage of this infrastructure and the number of supported end-user international scientific communities.

EGEE is proposed as a project funded by the European Union under contract IST-2003-508833