Workshop Abstracts

Topic 1: Towards a European Grid Infrastructure

- Dieter Kranzlmüller, Joh. Kepler University Linz, EGI Coordinator
  A European Vision and Plan for a Common Grid Infrastructure

The European Grid Initiative (EGI) represents an effort to establish a sustainable grid infrastructure in Europe. Driven by the needs and requirements of the research community, it is expected to enable the next leap in research infrastructures, thereby supporting collaborative scientific discoveries in the European Research Area.

The main foundation of EGI are the National Grid Initiatives (NGIs), which operate the grid infrastructures in each country. EGI will link existing NGIs and actively supports the setup and initiation of new NGIs. Within the current call of the EU Framework Programme 7, the EGI Preparation Team intends to submit a proposal for an EGI Design Study, which is currently being discussed with the NGI Representatives in Europe. The goal of this EGI Design Project itself is to evaluate use cases for the applicability of a coordinated effort, identify processes and mechanisms for establishing EGI, defining the structure of a corresponding body, and ultimately initiating the construction of EGI.

Paper: "The Future EuropeanGrid Infrastructure - Towards a Common Sustainable e-Infrastructure"
Vision Paper prepared for the EGI Workshop, Munich, 26-27 February, 2007 http://www.eu-egi.org/vision.pdf

- Steven Newhouse, Director OMII-UK
  User requirements for and concerns about a European e-Infrastructure

The plans for the emerging European Grid Infrastructure (EGI), built from a federation of National Grid Infrastructures (NGIs), will provide many managerial, operational and technical challenges. To ensure that the NGIs retain enough flexibility to meet the research needs and priorities of their national user communities it is essential that they are able to deploy the software of their choice provided that it implements the service interfaces required by the EGI.

Having service interfaces that conform to defined specifications is essential not only for interoperability with the EGI but also for the development of higher-level services, tools and applications that are essential in promoting adoption by applied end-users. A recent Survey of User Priorities for e-Infrastructure for Research (SUPER) undertaken in the UK amongst groups who had tried to use the e-Infrastructure highlighted several critical barriers to uptake by the broad research community.

The presentation will highlight some of the challenges in building the EGI from a software perspective but also how mechanisms to access and support the use of the EGI need to be considered in order to provide a usable infrastructure for the researcher.

- Kimmo Koski, CSC Finland and HET
  Results of the HPC in Europe Taskforce (HET)

During 2006 there has been intense work in Europe in order to increase European competitiveness in high-end computing. Various new activities have been started in addition to the existing European grid projects Deisa and EGEE. ESFRI has publisheda roadmap including plans 35 new major European infrastructures in which one entry is the European HPC Service. High-Performance Computing in Europe Taskforce (HET) was established in June 2006 with a target to draft a strategy for European HPC Ecosystem with a focus in petaflop computing.

Basic tool for modeling the European HPC ecosystem is the performance pyramid. One of the key arguments from HET work is to develop the different levels in the pyramid in a balanced way: enable sufficient top-class resources, but at the same time invest considerably in boosting the collaboration, scaling the software, building the competencies and developing the national/regional infrastructures to supportin creating a competitive and sustainable European HPC service.

HET completed the strategy work in January 2007. The recommendations of the HET taskforce are presented and their possible impact is discussed in the presentation. A FP7 project proposal PACE (Partnership for Advanced Computing in Europe) for the preparatory phase of HPC service including petaflop centers is introduced. Next probable steps with some scenarios and opinions are also presented.

- Kyriakos Baxevanidis, CEC
  European e-Infrastructure: current status, opportunities, challenges

In launching the partnership for growth and jobs as a new start for the Lisbon strategy, the 2005 European Council called knowledge and innovation the engines of sustainable growth and stated that it is essential to build a fully inclusive information society, based on the widespread use ofinformation and communication technologies (ICT) in public services, SMEs and households.

To get, however, the full benefits from ICT, EU Member States need more ambitious plans to exploit them, reveals the Commission's first annual progress report on i2010 – the digital economy part of the EU's revised Lisbon strategy for growth and jobs.

In this political and economic context, the European driven e-Infrastructure[1] provides the largest in scale and richest in terms of integrated technologies European ICT-based facility that enables researchers across Europe to face today's big challenges in their fields and that fosters also the emergence of a new generation of ICT-based infrastructures for the good of the whole European and world economy. The e-Infrastructure provides, in other words, a unique platform that boosts research intensity and that enhances innovation capacity across Europe. The benefit coming from public investment in the field is clear as according to an impact analysis study in the context of the preparations for the European 7th Framework Programme for Research (FP7), each 1€ of public R&D leads to almost the same amount or more of business R&D investment

Important challenges that are faced concerning the evolution of the e-Infrastructures in the coming years include the sustainability, interoperability and openness of the model; the alignment of the various policies (national, organisational etc) around the operations and the further development of the facility; the provision of training to users/operators/developers; and the efficient governance of the facility across Europe. All above challenges are

[1] Domain-independent ICT-based Research Infrastructures designed to support research; they integrate in a seamless way networks, computers, SW, data resources, experimental and training facilities to enable collaborative science and engineering.

important FP7-priorities calling all stakeholders to action. In this path the e-IRG is expected to play a key political role.

- Victor Alessandrini, DEISA Project Director
  Enabling cooperative extreme computing in Europe

DEISA's original motivation was to act as a vector of integration of High Performance Computing (HPC) resources at the continental scale. The DEISA services have been tailored to enable seamless access to, and high performance cooperative operation of, a distributed park of leading supercomputing platforms in Europe.

The DEISA services are deployed on top of a dedicated high speed network infrastructure connecting computing platforms, using selected middleware. Their primordial objective is enabling capability computing across remote computing platforms and data repositories. Workflows across different platforms, transparent remote I/O, large file transfers, are starting to operate without inducing performance bottlenecks that would invalidate high performance computing. These services will bloom as the number of scientific users accessing different computing platforms in Europe will grow.

After quickly reviewing the existing services and the service provisioning mode of the DEISA research Infrastructure, we will discuss why their persistency is mandatory to pave the way for new complementary initiatives in the area of HPC. Some comments will be advanced about the relevance of the DEISA environment for the efficient operation of future European supercomputers, and the current vision about the overall role of DEISA in the European HPC ecosystem will be discussed.

Topic 2: Sustainability for e-Infrastructures

- Costas Courcoubetis, Dept of Computer Science, Athens University of Economics and Business; and Richard Weber
  Sharing policies and resource provisioning in Grids

Grid technology is a virtualization technology. It allows the computing resources that are provided by various participants to be shared in a transparent way. By allowing resources that are owned by different organizations to be shared by others, the overall system efficiency increases through less idle time and better utilization of resources. This efficient multiplexing and use of economies of scale has been a main driving force in the development and acceptance of grid technology. It is widely believed that increasing the degree of sharing is the universal medicine by which to achieve economic efficiency. But is this really true? Are there any negative side effects of resource sharing that one should take into account to realize these multiplexing gains? Could grid participation actually reduce the benefits of the participants? As we will discuss in this talk, the answer is yes.

In order to understand the complete efficiency issue, one has to do a more complete assessment of the various parameters that affect the final system when grid technology is deployed. In particular, one has to appreciate that there are two interacting phases that influence the final outcome. The first is resource provisioning and the second is resource sharing. During resource provisioning, the grid participants decide on the amount of resources that each will contribute to a shared resource pool. During subsequent resource sharing, these resources are shared according to a policy specified by the administrator of the system. The key aspect that often eludes system designers is that, although once the amount of resources to be provisioned has been determined sharing then increases efficiency, the fact that the resources will be shared will influence participants' initial decisions regarding the amounts of resources they will contribute to the pool. More precisely, knowing that a large common resource pool will probably be available, a participant may decide to contribute a smaller amount of resources to the pool than is helpful, since he expects that most of the resources he might need will be provided by other participants. This

`free-riding' effect of resource sharing may have very negative consequences and even reduce the overall effect of grid at the equilibrium. It may also decrease the stability margin of the system and make the equilibria unstable.

As we discuss in this talk, the decision on how much resources to provision is greatly influenced by the resource sharing policy that will be deployed. The problem of choosing a policy which optimizes the efficiency of the system, assuming the amount of resources to be given may not be optimal overall, since it must take also into account its effect on the incentives during the resource provisioning phase. As we show in terms of a simple model, there can be an overall reduction of efficiency: a paradox in which, by building a common resource pool, the overall system efficiency is less than it is when building stand-alone systems for each participant. Initial work on stability also suggests that there may be serious stability issues if this policy is not appropriately chosen. However, we show that the right choice of sharing policy can greatly benefit the system and obtain the large efficiency gains that are anticipated by grid. Hence the incentives for resource provisioning play an important role and cannot be examined in isolation from resource scheduling and sharing. Designing optimal protocols for obtaining resources from the common resource pool is thus an important but difficult problem that must take into account the gaming aspects of the resource provisioning phase.

- Daniel J. Veit, Professor of Business Administration and Information Systems, E-Business and E-Government, Business School, University of Mannheim, Germany
  Grid Economy and Business Models; from Web to Grids

In the past few years substantial work has been done on protocols, middleware, and services for security, discovery, sharing, management etc., of computing, storage, data and other resources within dynamic, distributed computing systems. However, the economic aspect of allocating such resources has been underemphasized. In this contribution, the key concepts of integrating economic aspects into Grid resource allocation are proposed. Based on these, fundamental ideas for the integration of Grid services into business models for service providing companies and departments are outlined.

The latest and actual economic resource allocation methods for Grids from two large EU projects, CatNets and SORMA, are introduced. On the basis of these concepts, recommendations for the utilization of economic resource allocation schemes are outlined for the integration into European infrastructures.

Recommendation 1: Modeling of a two tier cost-oriented resource allocation scheme for E-Infrastructures based on a hybrid centralized/decentralized market mechanism for Grids and distributed architectures.
Recommendation 2: Introduction of a billing scheme for post-allocation resource billing in E-Infrastructures. Enabling of a cost-/profit-center concept for the utilization or computational resources.

- Neil Geddes, UK National Grid Service *Building a sustainable European Grid Service; lessons learned from a national experience*

I will review the progress towards building a Grid Infrastructure to support UK research and the challenges that remain in making this infrastructure sustainable. A particular challenge is to develop a shared intellectual ownership of the infrastructure which will allow it to grow organically. Different definitions of "infrastructure" and the often orthogonal demands of current stakeholders make achieving this shared vision challenging. For long term sustainability it is important that the current palette of activities grow collaboratively through a shared vision. Strong strategic commitment from, and coordination between, funding agencies and key user communities remains an important driver for the current infrastructure development."

- Prof Rob Procter, National Centre for e-Social Science, Manchester University
  Challenges for Sustainability: Perspectives and experiences from e-Social Science

- John Wood, CEO CCLRC and ESFRI Chair
  A European roadmap for new, large-scale Research Infrastructures and for a common e-Infrastructure

Research Infrastructures (RIs) are essential for developing top-class research activities, both basic and applied. Because of their ability to assemble a 'critical mass' of people and investment, they contribute to national, regional and European economic development. They also boost – through the development of skills, technologies and knowledge – the EU's chances of achieving its Lisbon Agenda, which is to ensure job creation and sustainable growth in what European leaders came to call the 'knowledge-based society'. Among the areas where RIs are most prominent in Europe are high-performance microscopes, laboratory clean rooms, very large telescopes, and specialised archives, libraries and databases. Such infrastructures could be 'single-sited', 'distributed' across multiple sites or, indeed, 'virtual' – where scientists can share data and carry out their work across the internet and through virtual office spaces.

Over the past two decades, the trend towards more integrated multidisciplinary research has rubbed off on pan-European research facilities which, instead of working exclusively on one scientific discipline or another, choose to provide services for a variety of scientific communities.

The process of creating an instrument for helping to identify those projects that are crucial for the scientific community in Europe has arrived to its first milestone. Europe's first ever Roadmap of large scale research facilities/research infrastructures, the ESFRI Roadmap (done in cooperation with e-IRG), has been published in autumn 2006.

Today, Europe RIs cover the full range of scientific and technological fields, from social sciences to astronomy, including genomics, energy and environmental sciences, nano-technologies and nano-sciences. What's more, RIs have ventured from the physical to the virtual world, as it now encompass high-performance communication networks, such as Géant, and new supercomputing 'Grid' technologies which take advantage of spare processing capacity in computers linked to each other all over the world.

_____

Géant's launch in 2001 was a hallmark event for European research, allowing over 3 000 research and education facilities in more than 30 countries to network.

Topic 3: Bridging the gap between academia and industry

- Ulf Dahlsten, Director, CEC; represented by Mrs. Lieve Bos
  Pre-commercial Procurement of innovation; A missing link in innovation in Europe

- Tony Hey, Corporate Vice President Microsoft, former Head of UK e-Science Programme
  From research infrastructures to industry participation; a long way to go ?

This talk will review the progress towards building a Global Research Infrastructure and look at the principal components of this infrastructure. It will then look at the issue of industry involvement in infrastructure standards and argue that the research community's priorities and industry are not yet well aligned. Arguably, companies like Amazon and Google, for example, are already running large Grids but see no commercial advantage in engaging with research community to talk about standards. Companies will not seriously collaborate in building cross-organization infrastructure components until they are convinced these represent real commercial opportunities. Until then, the global research infrastructure vision will remain classified as 'R&D'.

- Bob Jones, EGEE
  EGEE's vision and roadmap to involve industry

- Thierry Priol, INRIA, Scientific Coordinator of CoreGRID
  An e-Infrastructure for Grid Research and the Industry: the CoreGRID experience

- Michael Resch, Director HLRS, University of Stuttgart
  Collaboration with industry; from HPC to Grids

HLRS has worked with industry in HPC for at least 20 years now. While initially this was a system or project based approach HLRS has formalised the collaboration in 1996. At that time a company was set up to operate HPC systems and make them available to industry on a regular basis. Over time the focus of HPC in Stuttgart moved from operation of systems towards supplying services to users in general.

The collaboration with industry brought a number of benefits for HLRS including a constant income from the selling of CPU time. On the other hand industrial users brought with them a number of requirements that later on become know to the scientific community under the title of "Grid". A lot of what has been discussed for scientific grids in the last 5 years has been in the focus of industry for much longer. Hence HLRS has a very good understanding of what industrial requirements are and to which extent a purely scientific approach can fulfill these requirements. The talk will cover these aspects and will also draw the attention to some basic problems that divide the industrial world and the scientific world in HPC.