

Workshop report



Workshop organised and hosted by:

The General Secretariat for Research & Technology (GSRT),
The Greek Research & Technology Network (GRNET),
The Greek National Documentation Centre (EKT)
and
The European Commission

Document information

Identifier: eInfrastructures-Workshop-Report-V1.5
Preparation date: 24th July 2003
Status: Final
Author: Dr Mark Parsons, EPCC

Contents

1	EXECUTIVE SUMMARY	3
2	WORKSHOP AGENDA	4
3	INTRODUCTION	7
3.1	CONTEXT OF THE WORKSHOP	7
3.2	WORKSHOP OBJECTIVES	7
3.3	WORKSHOP AND REPORT STRUCTURE	8
4	SUMMARY OF PRESENTATIONS	9
	WELCOME	9
	<i>Evangelos Bouboukas - Welcome</i>	9
4.1	SESSION 1: NEW CHALLENGES FOR EUROPE ON EINFRASTRUCTURES	9
4.1.1	<i>Dimitris Demiozos – eInfrastructures and ERIA</i>	9
4.1.2	<i>Yannis Kalogirou – Broadband & eInfrastructures: a path to Regional Development and InfoSociety</i> ..	10
4.1.3	<i>Spyros Konidaris - Global eInfrastructures – The EU leading the way</i>	10
4.2	SESSION 2: EU PERSPECTIVES	12
4.2.1	<i>Mario Campolargo – The EU eInfrastructure initiative</i>	12
4.2.2	<i>Jean-Louis Picqué – The eInfrastructure Initiative</i>	13
4.2.3	<i>Kyriakos Baxevanidis – Towards a common European Networking & Grids infrastructure area – how can it work?</i>	14
4.3	SESSION 3: A EUROPEAN NETWORKING/NREN PERSPECTIVE	15
4.3.1	<i>Enzo Valente: Moving towards European Research Area: the two sides of the coin</i>	15
4.3.2	<i>Dany Vandromme: An eInfrastructure in Europe: an NREN perspective</i>	15
4.3.3	<i>Vasilis Maglaris: European NRENs and GTREN</i>	16
4.3.4	<i>Claire Milne – Research networks and the new European regulatory framework – competing public priorities?</i>	18
4.4	SESSION 4: NATIONAL/REGIONAL INITIATIVES	19
4.4.1	<i>Tony Hey – Building a European eInfrastructure: the view from the UK</i>	19
4.4.2	<i>Peter Kacsuk – How to build an inexpensive production Grid infrastructure</i>	20
4.4.3	<i>Manuel Delfino – The role of local and regional coordination</i>	21
4.4.4	<i>Mirco Mazzucato – The Grid infrastructure in Italy</i>	22
4.4.5	<i>Walter Hoogland – Perspective from the Netherlands: a bottom-up approach</i>	23
4.4.6	<i>Aleksander Kuszniir – CEGC: a view from the local, regional and European perspective</i>	24
4.4.7	<i>Marcel Kunze – Perspectives of Grids and e-Science in Germany</i>	24
4.4.8	<i>Lennart Johnson – Grid deployment and support – the NGC, EGSC and SweGrid initiatives</i>	25
4.5	SESSION 5: APPLICATION INITIATIVES	26
4.5.1	<i>Hans Falk Hoffmann – Grids and LHC: towards a first global Grid prototype</i>	26
4.5.2	<i>Fabrizio Gagliardi – Towards a common European market for computing and data management</i>	27
4.5.3	<i>Mark Parsons – The Grid: challenging HPC infrastructure provision in Europe</i>	28
5	PANEL DISCUSSION AND WRAP-UP	30
5.1	SUMMARY OF DISCUSSION	30
6	RECOMMENDATIONS	34
7	CLOSING	35

1 Executive summary

- Recent advances in Europe in the area of Research Infrastructures have resulted in the deployment of the GEANT high-speed research network to over 30 countries across Europe.
- At the same time a number of research initiatives in Europe and the World are creating pilot implementations of a model of shared use of computing and data resources across technological, administrative and national domains – the so-called Grid computing model.
- These developments are creating the expectation that Grid and networking technology is maturing quickly enough to support the emergence of a new infrastructure paradigm that in time will come to be seen as a commodity service.
- The creation of such an *eInfrastructure*, which will provide fully integrated communication and information processing services, is a key objective of the Research Infrastructures activity of the European Commission's Framework 6 programme.
- The full exploitation of a new innovative technological paradigm with such a broad scope and cross border relevance like the eInfrastructures concept can better (and in some cases only) happen when the appropriate administrative and policy decision mechanisms are put in place.
- The workshop described in this document was held in Athens in June 2003, under the auspices of the Greek presidency of the EU, in order to discuss the creation of the necessary administrative and policy decision mechanisms for the successful deployment of eInfrastructures within the European Research Area.
- The workshop was organised around a series of presentations from key EU actors in the fields of Networking, Grids and e-Science. Following these presentations a wide-ranging discussion was held which then endorsed a series of recommendations.
- Key recommendations arising from the meeting include:
 - The European Research Area should clearly be seen to embrace Innovation – articulated in the context of this meeting through the concept of ERIA.
 - An eInfrastructures Reflection Group, built from National Programme representatives, should be established and perhaps advise the governmental representatives who sit in existing committees.
 - The eInfrastructures Reflection Group should consider and communicate clear messages on eInfrastructure Policy issues to both the European Commission and existing eInfrastructure projects on policy matters.
 - It is clear that many countries are joining together into *regions* and this was presented as a powerful tool for cooperation.
 - The next steps for the Grid must be to move to *reliable, resilient, and robust* production quality middleware.
 - The idea of an Open Middleware Infrastructure Institute for Europe was broadly supported.
 - To build eInfrastructures we need to focus on middleware interoperability and the accompanying policy decisions required to make our software and operating paradigms interoperable in a global context.
- The meeting ended with a proposal for a future meeting to discuss progress on the many issues raised during the forthcoming Italian Presidency of the European Union.

2 Workshop agenda



The EU eInfrastructures Initiative

Towards integrated networking and Grid infrastructures for e-Science and beyond

Thursday 12th June 2003

Under the aegis of the Hellenic Presidency of the European Union

Workshop Chairs

Mario Campolargo EC Vasilis Maglaris GRNET

Local Host

Dimitris Deniozos GSRT Vasilis Maglaris GRNET
Evangelos Bouboukas EKT

Organizing Committee

Maria Christoula GSRT Kyriakos Baxevanidis EC
Jorge-A, Sanchez-P GRNET Fotis Karayannis GRNET
Stelios Sartzetakis GRNET Krystallia Drystella GRNET
Maria Koutrokoi EKT

Thursday, June 12, 2003				
National Hellenic Research Foundation, 48, Vas. Constantinou Avenue, Athens				
	08:30	30 min	Registration and Coffee	
0.1	09:00	10 min	Welcoming of participants and greetings	Evangelos Bouboukas, Director, National Documentation Centre
Session 1: New challenges for Europe on eInfrastructures				
1.1	09:10	15 min	eInfrastructures and ERIA	Dimitris Deniozos, General Secretary for Research and Technology, Ministry of Development, Greece
1.2	09:25	15 min	Broadband & eInfrastructures: a path to Regional Development and Information Society	Yannis Kalogirou, Secretary for the Information Society, Ministry of Economy and Finance, Greece
1.3	09:40	15 min	Global eInfrastructures – The EU leading the way	Spyros Konidaris, Advisor to the Director General, European Commission, - DG INFSO
Session 2: EU perspectives				
2.1	09:55	15 min	The EU eInfrastructure initiative	Mario Campolargo, Head of Unit, European Commission – DG INFSO

Workshop Report: The EU eInfrastructures Initiative

2.2	10:10	15 min	Policy aspects in support of the EU eInfrastructure Initiative	Jean-Louis Picqué, European Commission - DG Research
2.3	10:25	15 min	Towards a common European Networking and Grids infrastructure area - next challenges	Kyriakos Baxevanidis, Scientific Officer, European Commission - DG INFSO
Session 3: A European Networking/NREN perspective				
3.1	10:40	15 min	Moving Towards European Research Area: the two sides of the Coin	Fernando Liello, Chairman, European NREN Consortium Chairman
3.2	10:55	15 min	An eInfrastructure in Europe: A European NREN perspective	Dany Vandromme, Director, RENATER
3.3	11:10	15 min	European NRENs and GTREN	Vasilis Maglaris, Chairman, GRNET
3.4	11:25	15 min	Research networks and the new regulatory framework – competing public priorities?	Claire Milne, Antelope Consulting
Session 4: National/regional initiatives				
4.1	11:40	10 min	Building a European eInfrastructure: The urgent need for an Open Middleware	Tony Hey, Director, e-Science Core Programme, EPSRC
4.2	11:50	10 min	How to build an inexpensive production Grid infrastructure?	Peter Kacsuk, MTA SZTAKI
4.3	12:00	10 min	The role of local and regional coordination	Manuel Delfino, Director, Port d'Informació Científica (PIC)
4.4	12:10	10 min	The Grid Infrastructure in Italy	Mirco Mazzucato, Director of Research, INFN
4.5	12:20	10 min	The promotion of an eScience environment: a view from eSciencepark Amsterdam	Walter Hoogland, Dean Faculty of Science, University of Amsterdam
	12:30	60 min	Lunch break + Press conference	
4.6	13:30	10 min	CEGC: view from local, regional and European perspective	Aleksander Kuszniur, Deputy Director Cracow Academic Computer Center
4.7	13:40	10 min	Perspectives of Grids and e-Science in Germany	Marcel Kunze, Grid-Computing Kompetenzzentrum, Karlsruhe
4.8	13:50	10 min	Grid deployment and support – the NGC, EGSC and SweGrid initiatives	Lennart Johnsson, PDC, Royal Institute of Technology (KTH), Sweden
Session 5: Application initiatives				
5.1	14:00	15 min	Grids and LHC, towards a first global Grid Prototype	Hans Falk Hoffmann, Director, Technology Transfer & Scientific Computing, CERN
5.2	14:15	15 min	Towards a common European market for computing and data	Fabrizio Gagliardi, DataGrid General Manager, CERN

Workshop Report: The EU eInfrastructures Initiative

5.3	14:30	15 min	The Grid – Challenging HPC infrastructure provision in Europe	Mark Parsons, Commercial Director, EPCC and NeSC
	14:45	30 min	Coffee Break	
Session 6: Conclusions				
6.1	15:15	165 min	Open discussion and wrap-up	Round table: Spyros Konidakis, Mario Campolargo, Jean-Louis Picqué, Dany Vandromme, Vasilis Maglaris, Tony Hey, and Hans-Falk Hoffman
6.2	18:00	30 min	Grid technology showcase - Teravision- A Virtual Collaboration Environment technology	Electronic Visualization Lab-University of Chicago/GRNET

3 Introduction

3.1 Context of the workshop

Recent advances in Europe in the area of Research Infrastructures have resulted in the deployment of the GEANT high-speed research network to over 30 countries across Europe. This network is currently the fastest in the world. It builds on the experience gained over many years by DANTE and the NRENs across the Member States of the EU. In the context of GEANT, a policy committee consisting of representatives from all of the NRENs defines common rules of access to the network, common approaches for its rolling upgrade programme and coherent views of the upgrades proposed by NRENs.

At the same time a number of research initiatives in Europe and the World (building on the availability of high-speed networks and broadband connections, advanced virtual environments and new software technologies) are creating pilot implementations of a model of shared use of computing and data resources across technological, administrative and national domains – the so-called Grid computing model. This work is being supported and encouraged by the European Union and a number of national initiatives in a series of Grid middleware and applications development projects.

These developments are creating the expectation that Grid and networking technology is maturing quickly enough to support the emergence of a new infrastructure paradigm that in time will come to be seen as a commodity service. The creation of such an *eInfrastructure*, which will provide fully integrated communication and information processing services, is a key objective of the Research Infrastructures activity of the European Commission's Framework 6 programme. The acceptance and use of such an infrastructure has the potential to dramatically change the way in which people work and do business over the Internet and this is the reason why Grids are seen by many people today as the enabling technology for the next generation of science and business applications.

Further to these research efforts, experience has shown, particularly in the networking context, that the full exploitation of a new innovative technological paradigm with such a broad scope and cross border relevance like the eInfrastructures concept can better (and in some cases only) happen when the appropriate administrative and policy decision making mechanisms are put in place, in close alignment with the technological advances.

It was proposed in the context of this workshop to discuss the creation of an initiative on the policy level to discuss and formulate views on the above issues and to facilitate the creation of the necessary administrative and policy decision mechanisms for the successful deployment of eInfrastructures within the European Research Area.

3.2 Workshop objectives

The Workshop set itself the following objectives aimed at stimulating at a pan-European level an open discussion and to make recommendations for further actions to create a new EU policy initiative on eInfrastructures aimed at:

- Providing a framework for easier, faster and more cost-effective access to all types of information resources (networking, computing, data storage) distributed across Europe for all European researchers.
- Promoting the best practice implementation of such a model in the research area to facilitate and accelerate the subsequent commercial uptake of the new paradigm.

The more concrete objectives of the proposed meeting in Athens were to:

- Define the main shape of the new policy initiative in Europe - significant progress in this area may lead to the formulation of a new policy for approval by the Council of Ministers of the EU.
- Discuss the opportunity for the creation of a high-level committee of representatives from the administrative authorities of networking-, computing- and data-resources in Europe (similar to the policy committee of GEANT) to monitor and support on a policy advisory level the process of the creation and use of an eInfrastructure for Europe.

3.3 Workshop and report structure

The EU eInfrastructures Initiative workshop was arranged and organised under the aegis of the Hellenic Presidency of the European Union by The Hellenic General Secretariat for Research and Technology (GSRT), the European Commission and the Greek Research and Technology Network (GRNET) in collaboration with the Greek National Documentation Centre (EKT).

As is detailed in the workshop agenda, which is reproduced in Section 2, the meeting was organised around a series of presentations from key actors in the fields of Networking, Grids and e-Science. Following these presentations a wide-ranging discussion was held which then endorsed a series of recommendations.

Each of the presentations that was given is summarised in Section 4 in the same order as they were given on the day. Section 5 summarises the subsequent panel discussion and wrap-up while the recommendations arising from this discussion are listed in Section 6.

The meeting attracted an attendance of over 100 key European actors from the fields of Networking, Grids and e-Science.

4 Summary of presentations

Each of the presentations given at the meeting is summarised in turn. The summary is based on the presentation itself and notes taken during the talk. They are presented in chronological order and are grouped into the six themes given in Section 3.3.

Welcome

Evangelos Bouboukas - Welcome

Evangelos Bouboukas, Director, National Documentation Centre, opened the meeting and warmly welcomed all of the participants to it. He emphasized the importance of the meeting as an opportunity to discuss the future of access to networks and other infrastructures in the European context. The meeting was co-chaired by Mario Campolargo (EC) and Vasilis Maglaris (GRNET).

4.1 Session 1: New challenges for Europe on eInfrastructures

4.1.1 Dimitris Demiozos – eInfrastructures and ERIA

The European Research Area (ERA) is the central pillar of Framework Programme 6 (FP6) and beyond. It is a key component of the strategy to make the European Union (EU) the most competitive knowledge-based economy globally by the year 2010. The Greek presidency of the EU has stated the importance of innovation with regard to competitiveness – encapsulated in the idea of the European Innovation Area – and considers that the convergence of these two equal components should be represented as the European Research and Innovation Area (ERIA). The objectives of ERIA should be to: create an internal market in research, knowledge, researchers and technology to stimulate competition and better allocation of resources; to improve coordination of national research activities and policies; and to revisit the “subsidiarity principle” to understand how European thinking may influence national approaches – it is not enough to only reflect national priorities in EU strategies. The Greek presidency of the EU is committed to the ERIA process.

In the context of Research Infrastructures (RIs) and European competitiveness, which have been supported for the last decade through the Framework Programmes, the Competitiveness Council in Brussels agreed on the 3rd March 2003 that a high priority should be given to the Information Society in order to meet the objectives of eEurope identified in Lisbon and Barcelona. We must therefore debate the need for mechanisms: to jointly identify new research and technological challenges and how to respond to them in a rapid and effective manner; to increase the financial support by both the public and private sectors for RIs; to strengthen the role of RIs to lift obstacles to mobility and promote the integration of European scientists; to increase the contribution of national infrastructures to ERA; and to strengthen support for the successful deployment of Grid enabled eInfrastructures across Europe making full use of the GEANT pan-European research network.

The objectives of the eInfrastructure initiative are fully in line with the objectives of ERIA. They can be seen as providing a framework for: easier, faster and more cost-effective access to all researchers in Europe; allowing seamless access to information resources distributed across Europe; strengthening equal opportunities for all European scientists; and as a means to spread the benefits of “*big science*” to less advanced, remotely-located regions throughout Europe. These infrastructures should not be seen only as instruments for advanced science. They involve the development and use of many advanced technologies and many innovative solutions. It is crucial to promote best practice in RIs rapidly and to accelerate commercial uptake.

Greece is catching up in developing its knowledge based economy and is showing strong overall trends in improved innovation performance. GSRT funds technological infrastructure in Academic

and Research Institutes. These Research Institutes are ready to use and provide Grid resources connected to GEANT through GRNET, the Greek Research and Technology Network, which operates the national network at speeds up to 2.5Gbps. Greece is also supporting the extension of ERIA to the Balkans and Mediterranean countries.

4.1.2 Yannis Kalogirou – Broadband & eInfrastructures: a path to Regional Development and InfoSociety

The Greek Government believes that eInfrastructures represent the future of the Information Society. Greece itself is an excellent example of the importance of this statement. The emergence of a ubiquitous Information Society in Greece is a prerequisite for the convergence of the Greek economy with the EU average. A key point is that the Information Society is not just for top universities and large companies; it should represent society at large.

The eEurope Action Plan has set a number of challenging objectives for 2005 that include: broadband connections for all public administrations, schools, universities, museums and libraries; widespread availability and use of broadband networks throughout the EU; and the reduction of barriers to broadband deployment. Where applicable the eEurope plan supports the use of structural funds to achieve these objectives in less favoured areas. Broadband networking is central to the eEurope Action Plan and in many ways it can be seen as the “railway network” or “electricity power grid” of this century. In any country, broadband deployment requires a clear strategy and political commitment and this is magnified in Greece due to its rural nature and all of the challenges that this brings with it.

In the context of the Information Society there is a need for equal opportunities to make use of resources, independent of location or affiliation of the user. This is just as important for nations building their knowledge economy as it is for mature knowledge economy countries. The Greek Government supports the creation of a pan-European distributed environment for the provision of computing and storage resources to support scientists from across Europe. This statement has brought with it many challenges: Grid technologies are not yet mature enough to apply at a national level and furthermore the Greek market is reluctant to invest in these new technologies; there was also no national body for Grids in Greece (HellasGrid has now been created to meet this need and also to represent Greece in pan-European efforts like EGEE); and there was no separate funding for Grid projects (funding has had to be found from OPIS – the operational programme for the Information Society in Greece).

The Greek Government has taken specific actions to foster the deployment of broadband in Greece. This has involved the funding of broadband infrastructure and services deployment through OPIS. In the context of these actions, there has been a specific need to consider regulatory issues and wholesale prices in a deregulated communications environment. This environment has been structured in a way that encourages the development of a competitive broadband infrastructure for Greece.

To date, Greece has taken the lead in establishing eInfrastructures in South-East Europe through the work of GRNET in SEEREN. There is still much to do. For instance, Grid technologies are not yet mature enough to apply in the business domain. We must focus on defining policies for resource sharing, accounting, trust and security so that in future Grids may serve both the science and business communities. Only then will we be able to say that Grids are the “railways” of the 21st Century.

4.1.3 Spyros Konidaris - Global eInfrastructures – The EU leading the way

Greece has a long history of gathering and sharing knowledge. This Greco-Roman approach of creating, sharing and propagating knowledge in an environment open to everyone is embedded in Western civilisation and points to what global eInfrastructures can offer civilisation. In contrast,

consider the Egyptian model of privileged elites where most of the knowledge was restricted within the walls of the Pharaohs temples and much of which is as a result now lost.

It is instructive to consider the major milestones we have witnessed over the past 3-4 years. On the 31st October 2000, GEANT was legally born. GEANT reflected the NRENs desire to be more ambitious and have a clear vision, articulated in their charter, for European networking. As a result we have witnessed a meteoric rise in bandwidth from 155Mbps to 10Gbps. Although GEANT represented a large risk it is clear this has paid off – particularly in comparison to the Internet 2 infrastructure in the US which is currently only providing 2.5Gbps.

Shortly before this, on the 20th June 2000, the first Grid Workshop was held at the European Commission. This workshop clearly resonated both with the European scientific community and with the EC. Very shortly thereafter, funding was made available and serious Grid projects were launched in Europe – DataGrid, EuroGrid etc. During this time the speaker was in the US working as an EU Fellow. In talks and meetings two messages were repeatedly stated. Firstly, Europe is leading the world in its deployment of the GEANT infrastructure and secondly, in doing this we are not thinking regionally but globally. In this regard the US had to acknowledge the reality that, for once, it lagged Europe.

To broaden this discussion we must consider globalisation and cooperation. Today there are both positive and negative connotations associated with globalisation. On the positive side it is clear that communications networks (predominately telecommunications networks) have been the drivers behind globalisation. Our challenge now, in order to mitigate the negative effects of globalisation, is how to find a mechanism of plausible fair governance in order to make the most of it. As the performance of networks increases and as they become all pervasive, the process of globalisation will deepen. In this context the process of evolution of eInfrastructures is worth observing since they represent the spearhead of the network evolution with an impact much greater than just research and education.

Considering cooperation, the history of mankind has been determined by the equilibrium of two opposite forces: cooperation and competition. Both must exist, but we can now see clearly that the cooperation model will always win: for example consider the monotonic growth in the size of societal granularity from tribe, to village, to city, to state, ... to the world. This growth has been catalysed through the willingness of mankind to cooperate. As our problems become global the need to extend and deepen global cooperation at the expense of sterile antagonism becomes imperative. The European Union has been at the forefront of advancing this cooperative model for the past 50 years. This gives us the opportunity and the mandate to be at the forefront of the development of knowledge through global cooperation – this can be called the *Ecumenical Network of Knowledge*.

eInfrastructures bring together the resources of powerful new network infrastructures with potent new tools – such as Grid technology – to make knowledge resources accessible on demand and under agreed rules of conduct to all. In the US the phrase “*democratisation of knowledge*” is used to describe similar ideas. In Europe we often call this “*e-Science*”. In this context we must pursue a model which is not just regional but global. Although eInfrastructures are presently destined for an elite – the scientific research community – as they evolve and mature we will see their takeup first by business and professional users and eventually by ordinary citizens. A new end-user is being created, empowered and free to unleash unlimited human creativity.

In conclusion, the NREN community in Europe has assumed in recent years an acknowledged world leadership in the deployment of the most advanced network infrastructures. It has also become a major player in the development and deployment of Grids. By combining these two elements in the concept of eInfrastructures, and building on Europe’s cooperative skills we can set

Europe on the path to global leadership in this area. The Grid community has been afforded the opportunity to articulate a key vision – we must challenge our politicians to understand and support that vision.

Athens has always been a place that inspires new ideas. The location of this event is a happy coincidence – we should take advantage of the ghosts around us – Aristotle, Plato, Socrates ... to guide our thoughts and decisions.

4.2 Session 2: EU perspectives

Vasilis Maglaris introduced the second session and described how it focussed on presentations from three European Commission officials charged with developing the eInfrastructures concept.

4.2.1 Mario Campolargo – The EU eInfrastructure initiative

Many recent talks have concluded by saying that: the creation of virtual research organizations are key to rationalizing investments in e-Science; the EC has established an ambitious plan to implement the objectives of FP6 with regard to GEANT and Grids; and with this work the EC expects to deploy an eInfrastructure for all research communities thus contributing to the accelerated development of the European Research Area (ERA). However, they have also said that eInfrastructures require more than just technology – they also require policies and that the EC is ready to cooperate with Member States to devise and implement them. This should be the starting point for this meeting on eInfrastructures.

The highly successful deployment of GEANT and the early Grid pilot projects are clearly moving us towards eInfrastructures. GEANT is now IPv6 enabled and the Grid middleware is becoming more robust. As we tackle more and more of these problems we will find that the technological issues will gradually disappear and we will be left with an infrastructure for users that simply works. It is right that in these early stages the EC should focus on the scientific community and on eScience. But in the future this work must spread to eBusiness, eHealth, eLearning ... and EU citizens at large. Our goal must be to create an inclusive global knowledge infrastructure.

In terms of ERA it is clear that the eInfrastructure initiative must be one of its cornerstones, spearheading and expanding as it does the ideas of an eInfrastructure for Europe as described in the eEurope Action Plan. However, it goes beyond this by integrating national infrastructures, acting as a powerful instrument for international cooperation and contributing to policies such as cohesion, cooperation, standards, industrial competitiveness etc. We are already seeing direct examples of this in, for instance, the e-VLBI work, the infrastructure for the LHC at CERN, the HealthGrid applications and the early adoption of Grid solutions by some industries.

eInfrastructures need to serve both “normal” user communities and demanding user communities in a dynamic way. They must challenge the technologies employed by the research networks and push the developers of Grid middleware towards stable, robust solutions. We must address several levels of challenges: technology – middleware; organizational – virtual organisations; and policies. We need to articulate the policies we need at both a national and European (global) level and these policies should tackle issues such as: access to resources, geographical coverage, rationalisation of investments etc.

Developing these concepts will require the use of several instruments at both an EU and national level. In the EU context, the most powerful instrument is the Integrated Infrastructure Initiative (I3) supported by Specific Support Actions and Coordination Actions. We must also make use of national initiatives, regional/structural funds, ERANnet-like initiatives, and mobilising initiatives such as eEurope. In the context of national initiatives the idea of a steering group for national initiatives has been proposed. This steering group could: exchange information on the various initiatives, reflect on the challenges raised by a Europe-wide infrastructure, promote the adoption of

long term common strategies, policies and practices; initiate workshops; broaden user communities; trigger white papers; and give input to future National and EU workplans – FP6/7. We obviously must discuss the formalisation of such a body and this meeting is intended to act as a starting point for this.

In conclusion, eInfrastructures are a very ambitious concept that deserves a correspondingly ambitious approach. The current Grid testbeds have created high expectations that need to be met. The initial response to these ideas has been very positive and we have a real opportunity to lead worldwide – so long as we can move quickly.

4.2.2 Jean-Louis Picqué – The eInfrastructure Initiative

JL Picqué represented the EC's DG-RESEARCH at the meeting. The focus of their work is on supporting and building scientific communities across Europe. As such, Commissioner Philippe Busquin is strongly supportive of the eInfrastructures Initiative. Commenting on the earlier presentation from Dimitris Demiozos it was stressed that innovation has always been key to the European Research Area (ERA).

The main rationale behind ERA is to focus on the fragmentation of the European research landscape and to try and improve this situation within the context of the EU. ERA was originally proposed in January 2000 and has gathered broad support in political and scientific circles. Its implementation is ongoing and its main thrust is for open coordination of activities across the EU. eInfrastructures are an essential tool for the construction of ERA; they have the potential to connect more than 3000 research centres across Europe and give access to enabling infrastructures to all European Scientists regardless of their location. The expected outcome will be a structuring of scientific communities in the European context in various disciplines.

The overall budget for the Framework programmes has increased markedly over time. The total FP6 budget is €17.5 billion and this is shared between three major activities: integrating European research; structuring ERA; and strengthening the foundations of ERA. The largest of these activities being the goal of integrating European research with around 82% of the proposed budget allocated to it. In FP6 the Research Infrastructures action has an increased budget compared to FP5 of €655 million, which includes €200 million for GEANT and Grids. Its main objectives are to provide access to infrastructures irrespective of their location in Europe and to promote the optimum development of new and enhanced infrastructures. It is open to all fields of science and technology. There are a number of instruments associated with this action which include, for existing infrastructures, Integrated Infrastructure Initiatives, Transnational Access, and Communication Network Development. For new infrastructures the instruments include Design Studies and Construction of New Infrastructures.

In the context of this meeting, the Research Infrastructures action focuses on the Communication Network Development Scheme that is implemented by DG INFSO in conjunction with the IST priority thematic area. It covers high capacity, high-speed communication networks (GEANT) and high performance Grids and test-beds.

An FP6 Coordination Group has recently been set up, and includes DG RESEARCH and DG INFSO, to coordinate efforts between priority thematic areas and the Research Infrastructures action. The aim is to explain the strategy and actions on Grids and GEANT and to take into account the emerging needs from the user communities.

A little over a year ago, the European Strategy Forum on Research Infrastructures was set up by Member States to support a coherent and strategy-led approach to policy making in the context of European Research Infrastructures and to facilitate multilateral initiatives for the development of Research Infrastructures, in particular focussing on acting as an incubator for “variable geometry” arrangements. The EC provides support to this informal group of high-level national representatives

and five meetings have been held since April 2002. At the meeting on the 28th April 2003, the French delegation proposed the establishment of a Working Group on High Performance Computing & Networking. This could be used to identify science needs and to propose how to coordinate national Grid initiatives.

4.2.3 Kyriakos Baxevanidis – Towards a common European Networking & Grids infrastructure area – *how can it work?*

In 2003 the MIT Technology Review identified Grids as one of the “ten technologies that will change the world”. The advantages of the Grid approach for science and business are clear. Grids will transform the IT landscape from discrete infrastructure components to a distributed information processing model where people share and do not necessarily own IT-resources. Organisations can therefore focus on their business objectives – be they scientific or commercial – rather than on the management and maintenance of underused (in many cases) IT-infrastructures. Our aim must be to construct a “one stop shop” service for users providing them with access to IT-resources, which meet their needs, and thereby transform Grids into a public utility.

Of course, technical and process developments will drive this transformation. By focussing on solving the technical challenges of security, quality of service, and ease of use while understanding the central business needs of users and also where Grids can bring immediate benefit, we can move towards this vision. As we do this we will see the price/performance ratio of IT installations decrease as the homogeneity of policies for accessing and using these resources are better understood. Tackling these non-technical barriers – the need for global agreements and policies to enable global use – are key to the future development of eInfrastructures.

Platform Computing has recently published a survey focussing on the non-technical barriers to the widespread uptake of Grids. The results make interesting reading. As one major EDA chip manufacturer says in the report: “If we move to a global Grid, we need agreement on a global infrastructure ... We will be managing a cultural change; people will need to broaden the scope of their thinking”. Likewise a global auto manufacturer comments that: “When you try to build a Grid and you have to do it within a company, you have to set policies and guidelines and everyone has to agree to give up their own resources in a shared pool. A global infrastructure causes global problems”. In the Platform Computing survey a startling 89% of organisations identified organisational politics as a barrier to implementing Grid solutions in their organisations. The key conclusions from the report are that these non-technical policy aspects of Grids are significant barriers to their implementation – people in general do not have a resource sharing attitude. Moreover, very different policies for accessing resources across institutions, application domains and national boundaries in Europe exist. The harmonisation of such policies at all levels is therefore a major challenge.

The implementation of GEANT has taught us many things in this context. Tackling pan-European connectivity resource sharing policy aspects at a European level has resulted in the world’s fastest research network that provides affordable access to all researchers. This has been achieved through fully-fledged operational support and a policy committee to resolve policy issues. We have learnt that interconnecting *people* matters more than interconnecting *machines* when trying to meet our goals.

To meet the policy challenges created by eInfrastructures, we need to create structures and mechanisms to harmonise IT-resource access and use policies across Europe for e-Science and beyond. We should formulate an eInfrastructure policy framework and use it to establish appropriate administrative, operational and policy support schemes for IT-resource sharing at all levels. We should consider setting up a high-level expert group to monitor this process and provide

advice. We must ensure that all interests and groups are sufficiently represented and consider the allocation of EU resources to catalyse the process.

One in twelve citizens of this planet is a member of the largest common market in the world – the European Union. Can we afford not to establish common market structures for the use of our IT-resources?

4.3 Session 3: A European networking/NREN perspective

4.3.1 Enzo Valente: Moving towards European Research Area: the two sides of the coin

Enzo Valente presented this talk in place of Fernando Liello who sent his apologies.

In the context of networking and ERA, Europe has filled the gap with the GEANT network. This network is the most advanced in the world today with a core bandwidth of 10Gbps. It has exploited telecommunications liberalisation in Europe and built on the rich experience of the NRENs and the essential support of the national and European funding bodies. GEANT is directly in line with the concepts of subsidiarity and complementarity.

GEANT and the NRENs are mutually dependent on each other for their success – without the NRENs GEANT would be useless and vice versa in the European context. Taking this approach has provided a more complex but more flexible architecture able to meet the end-to-end challenge of providing connectivity across the continent for all research users. GEANT has two foci: it provides a network for research – based on advanced, transparent worldwide services – and also undertakes research on networking – based on a quickly evolving, segregated infrastructure for “risky” activities – without neglecting the overall need for sound operation.

From the point of view of GEANT and the Grid, the provision of services is key. Bandwidth for the sake of bandwidth is useless and researchers must be supported with the most advanced services achievable. In FP5 the successful and fruitful experiences with both the DataGrid and DataTag Grid projects have been very important.

In networking, global collaborations are clearly important when focussing on research excellence. In terms of FP6, the new countries that are acceding to the EU are putting pressure on the GEANT model. These countries are pushing the technology envelope forward and may well use dark fibre for instance. Our challenge is to fight the divide created by different telecommunication markets in Europe and work with the diversified procurement strategies evident across Europe to provide the best possible service to our user. In particular, there is a great need for close coordination between the Grid Research Infrastructures community and the networking community throughout Europe.

However, solving the connectivity issues in Europe is not enough. Research networks must be global as well. One issue is that international initiatives are not specifically included in the main stream of EU support. There are ongoing projects to connect emerging regional networks to GEANT, for instance in the Mediterranean, Latin America and Pacific Rim. These activities should be further strengthened and supported in FP6.

4.3.2 Dany Vandromme: An eInfrastructure in Europe: an NREN perspective

In order to understand the formation of eInfrastructures it is instructive to consider the history of networking infrastructures in Europe from the point of view of an NREN – in this case the French NREN, RENATER.

Like many NRENs across Europe, RENATER has developed in parallel with the European interconnection network. From the early 1990's, RENATER was an IP service established between a series of PoPs provided by the national operator – in many respects it was a black-box service. At the same time telecommunication companies provided a European service through a series of

packaged service – X25 with IXI, ATM with JAMES and IP with EUROPANET. These services provided a maximum of 2Mbps connectivity.

From 1996, France Telecom had to provide the IP service on a dedicated ATM infrastructure to meet the growing needs of users (to fill the gap between 2Mbps and 34/45Mbps connectivity) and to allow monitoring from the user edge. At the same time, TEN-34 was started and built on half-circuits provided by monopoly telecommunications operators on a very ad hoc basis. Bandwidth provided was around 10Mbps.

In 1999 RENATER was set-up as a major procurement action where circuits, equipment, PoP hosting and network management were sought. The outcome of this procurement was that France Telecom retained most of the circuit provision but lost the network management. Equipment was acquired directly by RENATER and PoPs were installed in Universities and Research Institutes around France. Similarly, TEN-155 was organised as a major procurement action by DANTE on behalf of the NRENs across Europe. A single operator won the circuit provision and the ATM layer while an NREN network operations centre managed the IP service. Equipment was acquired separately by DANTE.

From 2002, RENATER-3 has been built on WDM/SDH circuits. France Telecom has lost almost all of the circuit provision and the network is totally under the control of RENATER. By 2002, GEANT was also operational with eight different connectivity suppliers. Equipment was procured separately and the network management service was outsourced to a specialised company. The network is totally under the control of DANTE on behalf of the NRENs.

In the future we must be careful not to take over all of the tasks of telecommunications providers – they should be our partners – although we realise that research networks provision can never be “off the shelves” because they must remain innovative and at the leading edge. We need to work with the telecommunication operators to convince them to provide raw capacity at the lowest rates possible but leaving much of the mastering of the technology in the hands of the NRENs. Moving from SDH to WDM, from WDM to lit fibres and eventually dark fibres instead of lit fibres when this is feasible.

The key challenge is not to be destructive to the telecommunications providers but rather to provide them with incentives to provide new services. We must coordinate disparate actions, for instance the xx-Light initiatives that promote lambdas rather than usage and account for differentiated economical and regulatory contexts to harmonise the European network and reduce the digital divide.

TEN-155 and latterly GEANT have greatly improved the provision of pan-European networking but they have also increased the gap between European nations. Today for instance, the countries of South East Europe are amongst the most expensive for GEANT. We need to understand how to handle this challenge for the greatest benefit of the entire network – a challenge which has been confirmed by the EUMEDCONNECT tendering process.

Finally, it is clear that it is much more beneficial to cooperate with telecommunications companies and work with them to solve challenges. This is particularly important where little or no market exists due to the remoteness of some territories. In this regard France Telecom deserve thank from RENATER in the context of connectivity to remote French associated territories. It is always better to have one operator than none.

4.3.3 Vasilis Maglaris: European NRENs and GTREN

Over the last decade the European NRENs have adopted a “business model” characterized by the following traits.

Firstly, in the vast majority of European nations, a single state-controlled advanced infrastructure serves all Universities and Research Centers networking needs. These, apart from pure “research” electronic communications (between or among researchers), may in many cases include “commodity” traffic, i.e. traffic that has a source or a sink in the Research & Education community, while the other end is the global Internet. Transient “commercial traffic” defined as connections using the NREN as a “via” structure to serve two commercial entities is not compatible with the current regulatory and financial organization of GEANT and NRENs. This is clearly stated in the written agreement between NRENs and end-users, referred to as the *Acceptable Usage Policy* (AUP). Thus, NRENs do not compete with the ISP community and do not distort electronic communications markets in a highly competitive environment. It is worth mentioning that only 3-5 NRENs applied for a general license (authorization) to provide public electronic communications services, out of more than 30 belonging to TERENA (Trans-European Research & Education Networking Association). In Europe, the NREN community maintains the academic & research networking tradition that was the driving force for the *ARPAnet – NSFnet – Internet* early stages in the USA. The successors of the Academic – Research network in the US (*Internet2* initiative - *Abilene*, *vBNS+*, *ESNET* etc.) may be technologically advanced platforms for cutting edge applications (e.g. collaborative virtual environments with tele - immersing experiments, virtual distributed orchestras etc.) but are restricted to a small number of advanced users. For example, *Abilene* serves less than 200 Universities and Research Foundations (the *University Corporation for the Advancement of Internet Development – UCAID*); even within UCAID, *Abilene* serves a small minority of users (advanced *eScience* experiments), while the majority is being served by commercial ISPs. Few members of the US Academic community take advantage of the largely under-utilized *Abilene* resources, sometimes not even knowing of the option to use it. On the other hand, NRENs in Europe and their Pan-European gigabit interconnection GEANT, serve more than 3000 Institutions, half of which are using it as their sole gateway to the global Internet via a service provided by DANTE.

Secondly, the European Research & Education networking model evolved into a three-tier architecture: The campus LAN, the national MAN – WAN (the NREN) and the federal gigabit interconnection GEANT. All three tiers enable end-users to communicate with gigabit speeds as if the campus LAN is extended into the whole European Research Area. Apart from providing connectivity to researchers and the educational community, the three-tier structure may arrange for the provision of *Virtual Private Network* (VPN) resources to e-Science projects (e.g. GRIDs) on request, possibly with end-to-end Quality of Service guarantees (jitter, speed, security etc.) The strictly hierarchical structure of Research & Academic networking in Europe may exhibit scale economies in the provision and management of user, national and Trans-European resources, but may suffer from rigidity to follow the overall Internet paradigm, which is based on peering and neutral interconnection facilities (GigaPoP telehousing). The latter becomes evident when planning the introduction of user-empowered infrastructures in Research & Education Networking such as dark fiber ownership, condominium sharing (the Canadian business model) and long-term IRU leasing. Note, that at present GEANT is based on IP/MPLS provision over DWDM “lambdas” (or over SDH circuits in cases where DWDM circuits are not available), on short-term leases from international electronic communications operators. The three-tier hierarchical model does not encourage NREN clustering at regional levels; this may introduce a fourth level in the hierarchy or may eventually render the Trans-European level (GEANT) obsolete and replaced by peering arrangements.

Thirdly, the hierarchical model is interpreted by some GRID end-users as a nuisance, introducing complicated capacity management schemes (involving NOCs of campuses, NRENs and GEANT) to set-up high speed end-to-end connections that in some cases could be provided by a direct “lambda” circuit, bypassing LANs, NRENs and GEANT. Nevertheless, the hierarchical (federal) model has

been up to now a great enabler for Universal Service Provision and a means to bridge the digital divide across Europe. It is interesting to note that popular Pan-European multidisciplinary GRIDs are built or planned along the tier model, in fact imitating the NREN – GEANT paradigm. A reason may be the scale economy and organizational ease that this model achieves in managing vast shared computing and storage resources and the need for a strict trust schema, based on the tree concept of “root certification authority.”

Finally, the European “federal” Research & Education Networking platform attracted global interest as it unified thousands of advanced European researchers into a critical mass comparable or superior to US, Canadian and Japanese networked communities. Thus, the successor to GEANT is expected to be the driving force in the *Global Terabit Research & Education Network – GTREN*. As a first step, European NRENs (together with their non-for-profit organizations *DANTE* and *TERENA*) are tying together the European Research Area (GEANT) including South-East Europe (*SEEREN* initiative), North America (gigabit Transatlantic connections to *Internet2* and *Canarie*), South American (*@lice* initiative), Mediterranean countries (*EUMED-CONNECT* initiative), links to NRENs in the Russian Federation, Ukraine, Asia – Pacific (*TEIN* initiative) etc.

It is expected that *GEANT+* (the future GEANT upgrade) will continue to drive networking technology to its limits (e.g. optical switching, terabit capacities) and will help establish the European researcher as an ever-growing user of world-wide distributed *eScience* applications. This is exactly the driving force in deploying *GTREN* from a European perspective.

Planning of *GEANT+* will have to successfully resolve its biggest challenge, i.e to convince e-Science end-users of its capability in providing QoS enabled VPNs (at levels 1, 2 and 3) in a seamless, transparent mode to the user. Otherwise research and academic users will eventually drop-out from the established three-tier hierarchy in favor of direct connectivity solutions (via telco and/or owned optical links). As we very well know it is a jungle out-there, that researchers of the extended *European Research Area (ERA)* were able to overcome so-far thanks to the orderly, universally provided hierarchical GEANT - NREN infrastructure.

4.3.4 Claire Milne – Research networks and the new European regulatory framework – competing public priorities?

Telecommunications regulation in Europe is entering a new era and this is an important moment to consider the effect of these new regulations in terms of the NRENs. The SERENATE project has been working to assess the main implications for NRENs of the emerging regulatory situation particularly with reference to new ownership models, market liberalization and the rules for running networks and providing services. It should of course be noted that these rules exclude NREN-specific rules as described in their own statutes.

We are in the middle of a sea change in Europe. From July 2003, telecommunications regulation will cover all electronic communication, which clearly reflects the outcome of convergence in this area. Content regulation is excluded from these regulations and will be dealt with separately. What will be regulated in future are *services* not telecommunications *provision*. This means that NRENs may possibly come under the regulations because they receive remuneration for the provision of services. In theory all EU telecommunications markets have been fully liberalised since 1998 and the status of this is tracked by the EC’s annual implementation reports. The new regulatory package was approved in April 2002 and must be implemented in all states by July 2003. However, many countries are not yet ready for these new regulations and NRENs have an opportunity to influence national law in their favour. The Accession Countries must adopt these regulations *acquis communautaire* by their date of joining. The idea of a European Regulator was floated in 1999 but this has now been dropped.

The main points of the new regime include: the creation of the freest possible market consistent with adequate consumer protection; continuing the basic principles of regulation; abolishing licensing through the establishment of general authorisations for electronic communications service provision; and a market analysis procedure that must justify additional *ex ante* regulation to curb abuse of significant market power – aimed mainly at former incumbent operators. The implications for NRENs are generally positive. They will benefit indirectly from lower prices, increased choice and quality but the changes may bring some NRENs directly under the regulations and this may open up some issues.

Interconnection will now become a special case of access and is defined as “the physical and logical linking of networks to enable users of both networks to communicate with each other”. Public communications network providers are defined as providing wholly or mainly publicly available electronic communication services. They must negotiate their own access and interconnection contracts but the regulator may intervene when required – particularly with respect to operators who still exercise significant market power. NRENs are not generally classed as public communications network providers because they serve a closed community. There have been some worries from ISPs in this context with regard to unfair competition and this is acknowledged as a hard problem particularly as the number of users served by NRENs expands to schools, homes etc.

NRENs are funded for the public good and help close “digital divides” between and within countries. It is in the public interest for NRENs to get the best possible terms for interconnection and access even if they are not formally classed as public communications network providers. In this context public-private partnerships may be worth exploring for maximising the rapid provision of advanced infrastructures especially to less favoured areas.

4.4 Session 4: National/regional initiatives

4.4.1 Tony Hey – Building a European eInfrastructure: the view from the UK

The UK e-Science Programme is the largest national Grid initiative in Europe and one of the largest in the world. The funding has been split into two phases. In Phase 1 (2001 – 2004), application projects totalling £74 million of funding and a core middleware programme totalling £35 million of funding are presently underway. Phase 2 (2003 – 2006) has recently been confirmed and this brings with it funding of £96 million for application projects and a core middleware programme amounting to £41 million (although £25 million of this has still to be confirmed). By the end of 2006 the UK will have invested almost £250 million in e-Science and it is imperative we have a working infrastructure to show for it.

In Phase 1 the projects have largely been research and development projects and not production quality software engineering projects. The middleware that has been generated is not going to be easily deployed until more engineering effort is put in. Despite this, the programme has over 80 UK companies actively participating bringing a further £30 million of industrial contributions to the programme. These companies come from a variety of sectors including: engineering, pharmaceutical, petrochemical, IT, commerce and media.

To support these projects an e-Science Grid, focussed around ten University research centres and two government laboratories, has been established. Creating this Grid has proved very challenging and considerable experience has been gained in the practicalities of building a real heterogeneous Grid.

Over the past two years it has become clear that there is a need to develop an open source, open standard compliant middleware infrastructure which will integrate and federate with industrial solutions. Such work must have a software engineering as well as a research and development focus. The aim must be to produce robust, well-documented, re-usable software that is maintainable

and that can evolve to embrace emerging Grid standards. With this in mind, a major focus of Phase 2 of the UK e-Science Programme is the creation of an Open Middleware Infrastructure Institute. This will act as a repository for all of the UK-developed Open Source “e-Science/Cyber-infrastructure” middleware. It will also act as a document repository and involve itself in quality assurance and compliance testing for GGF/WS standards. A key role of this Institute will be to fund software engineering effort to bring “research project” middleware up to “production strength” – which we know can take an order of magnitude more effort to achieve compared to the original development. The Institute will also fund middleware development projects for identified “gaps” and work with US, EU and Asia Pacific projects. We intend that the work is supported by major IT companies.

In the EU context – should we consider a similar Open Middleware Infrastructure Institute for the middleware developed by the EU funded Grid projects? This would take a series of roles similar to those described above for the UK Institute.

In summary, there is a clear and urgent need for software engineering to develop a consistent eInfrastructure middleware stack. It is essential we have an Institute similar to the UK in the European context. This is a bold vision and calls for bold initiatives. We intend to invest £30 million in the UK; a similar amount of funding will be required in the EU. We believe this work, in both the UK and Europe, is vital to avoid a backlash from new users who find problems with what is currently “proof of concept” middleware. Unless we take coordinated action now we will not have a robust eInfrastructure for deployment by science and industry by 2007. As Tony Blair said in 2002, “[The Grid] intends to make access to computing power, scientific data repositories and experimental facilities as easy as the Web makes access to information”. Now is the time for action to meet this goal.

4.4.2 Peter Kacsuk – How to build an inexpensive production Grid infrastructure

Hungary has recently established the Hungarian Grid Competence Centre (MGKK) to lead an ambitious plan to coordinate Hungarian Grid efforts throughout higher education and research institutions. A key aim is to establish a cluster-based production Grid throughout Hungary’s higher education institutions. MGKK is a virtual organisation founded by four leading institutions – MTA SZTAKI, NIIFI, BME and ELTE. The organisation is focussing on two main projects: ClusterGrid – which aims to connect the Hungarian University clusters into a high-throughput Grid system and SuperGrid – which aims to connect the Hungarian Supercomputers into a high-performance Grid system.

The central goal of the ClusterGrid initiative is to connect 99 new clusters, which have been installed throughout the Hungarian University system, to form a production Grid. Each cluster consists of 20 PCs and a network server PC. During the daytime the components of the clusters are used for education. Overnight, the clusters are connected over the Hungarian Academic Network (2.5Gbps) to form a Grid. The total capacity of the Grid by the end of 2003 is expected to be 2079 PCs.

The basic concepts of the system have been to: keep the system as simple as possible; to use existing production quality network and Grid middleware components; to only develop missing components; and to utilise only one entry point for security reasons. The existing components that have been chosen are Condor (using its flocking mode for brokering) and VPN technology (solving the firewall issues with Condor). The new development undertaken has been to develop system boot software for the overnight Grid working mode. This software is designed to make switching between the different working modes as automated as possible. It runs continuously on the central Condor master and the local Condor masters.

Currently there are eight Hungarian Institutes involved (five are outside Budapest). Approximately 500 nodes are currently enabled for Grid operation. About 250 of these nodes are brought into the Grid each night and at weekends. The total number of nodes is rapidly increasing. They are seeing utilisation levels of 40% already – prior to the service being properly opened and expect this utilisation to increase when it is. The management of the project is arranged around a Technical Committee and a Steering Committee.

A number of further developments are foreseen to meet initial problems with the system. These include: a high level Grid programming environment is missing – this will be fixed by installing P-GRADE; there is no parallel checkpoint support in Condor – a joint development is underway to combine P-GRADE and Condor to provide this functionality; Condor job monitoring is not satisfactory – the GAMI software developed by MTA SZTAKI in the DataGrid and GridLab projects will be adapted; and the single entry point is a cause for concern due to overloading – there are plans to configure a separate entry-point machine for each site.

There is a clear “chicken and egg” problem in relation to production Grids: should a user community be established first or does this require an existing infrastructure? In the context of ClusterGrid, the infrastructure is being put in place first in order to be ready to meet the needs of users. Social issues have also played a role in terms of convincing the cluster owners to allow their machines to join the Grid. This has been solved by demonstrating the benefits of the Grid to the cluster owners through some early success stories.

The Hungarian ClusterGrid Initiative is demonstrating how to create an inexpensive production Grid system. Other countries are already showing considerable interest in this approach. It is hoped to be able to take the idea further and connect to other Grid systems such as DataGrid and, in due course, EGEE.

4.4.3 Manuel Delfino – The role of local and regional coordination

eInfrastructure enhancements will only succeed if we solve end-to-end issues at the technical, infrastructural, methodological and social/human levels. In this regard, local and regional coordination will be an essential ingredient for tackling these issues. The sophisticated structure of Europe is uniquely positioned to achieve this multi-level coordination. In the context of this talk a local area can be thought of as a city (eg. Barcelona), whilst a regional area can be thought of as an area of a country (eg. Catalunya).

Where appropriate, regions will provide the “flexible meso-level” between the national and local levels in order to catalyse the rapid and effective take-up of eInfrastructure technologies and methods of working. Regions should be seen as a complementary vehicle to the overall effort to achieve the vision of eEurope. Groups of regions acting in a coordinated fashion with a direct connection to actors in e-Science, Technology, Health and Industry will enhance the benefits of introducing eInfrastructures.

At the local level, the involvement of cities and metropolitan areas are essential to avoid “first kilometre problems” and to provide the ultimate link to citizens in terms of the technical, social, and organisational problems they may encounter. Typically the competences required at the local level will be different but complementary to those at the regional and higher levels.

There are some obvious examples of how these levels may be structured. In the health sector we already see that hospitals and health centres are organised at a combination of local and regional levels. In the Governmental sector there are an increasing number of agreements to present a unified “one-stop-shop” to the citizen for access to local, regional, national and EU services. From the point of view of consumers, they largely operate at a local and regional level and we must recognise this in the policies we set.

The GridPort concept encapsulates these ideas and is specifically designed to enhance the feedback between the growth of Grid infrastructures and the development and deployment of e-Science applications at the regional level. The Association of Regional GridPorts will catalyse the coordination and cooperation between individual GridPorts leading to the coordinated growth of e-Science communities in a scalable and timely way – hopefully mitigating the creation of an “e-Science bubble”. In particular they will focus on the compatible and cost-effective deployment and growth of Grid infrastructures and avoid the problem of unfulfilled expectations of new users. These ideas are being incubated currently by the Governments of Baden-Württemberg, Catalunya, Lombardia and Rhône-Alpes but will be open to all.

In summary, local and regional coordination will be an essential ingredient to solving end-to-end issues. With support from the Association of Regional GridPorts we believe we have a powerful tool to best utilise the uniquely sophisticated structure of Europe for the benefit of all science.

4.4.4 Mirco Mazzucato – The Grid infrastructure in Italy

Italy started to develop Grid technology and related infrastructures in the second half of 1999 through the INFN-Grid project. It was based on the realisation that modern science is moving to a new phase of global collaboration to improve efficiency, avoid duplication of effort, combine distributed expertise and build critical mass – all of this being encapsulated in the idea of e-Science. INFN represents 25 sites and realised the pressing need to integrate these sites using Grid technology as we move towards the development and deployment of a whole range of new scientific digital instruments and their associated data rates.

Grids clearly have much wider applicability than simply e-Science. Modern industries, business and Government are relying more and more on innovative solutions to problems and are increasingly basing their decisions on a cycle of problem modelling, simulation of various solutions, selection of best solution and realisation. To maintain and increase European competitiveness we must put in place the technologies which will allow us to quickly and easily assemble distributed teams that utilise distributed data and computing resources based on well understood collaborative methods – the central philosophy behind the creation of Virtual Organisations.

Towards the end of 1999, INFN decided, after careful evaluation, that Grid technology promised to deliver a key enabling solution to the problems faced by High Energy Physics and e-Science in general. The development of the components of a national Grid infrastructure for Italy has been actively pursued since then. This has been achieved through the active participation in Grid middleware development projects such as EU DataGrid and DataTag, promoting international collaborations (largely with US initiatives such as Globus, Condor, iVDGL, PPDG and GGF) to allow worldwide interoperability in projects such as GLUE, WorldGrid and MAGIC. As a result the Italian national Grid infrastructure is now a reality being developed by a series of coordinated national projects.

This development has been led by the INFN-Grid project, which was approved at the beginning of 2000. Initially this work was focussed on the preparation of the INFN LHC computing service but since then has become a more general solution. The project was complex – involving 20 Italian sites, ~100 people and a budget of ~€30 million. It represents a successful collaboration between physicists, software engineers, computer professionals, computer scientists and Italian industries. It has resulted in a reliable INFN Grid infrastructure involving all 20 sites and with a focus on support and general services.

The next stage has been to transform this work into an Italian Grid and this is being accomplished through the national FIRB Grid.it eInfrastructure project. This three-year project, which started in November 2002, has a total budget of €8.1 million and brings together, INFN, CNR, ASI and associated Universities. A wide range of sciences are now supported and this project has the

responsibility for creating a national Grid infrastructure and prototyping a national Grid Operation Service (GOS). In addition to this, and building from INFN-Grid and Grid.it, the Italian Grid for Business, Industry, Government, e-Science and Technology (IG-BIGEST) has also been established. Its aim is to promote the establishment of a general EU Grid infrastructure for e-Science integrating all of the available EU national infrastructures. This work is coordinated by INFN who see this project as a key enabler in meeting the vision of ERA.

Current challenges facing the creation of an EU eInfrastructure include: delivery, robustness and stability of middleware, managerial and administrative structures, and policies for resource sharing, virtual organisations, security, and accounting etc. It is hoped that the proposed EGEE project will tackle many of these issues and it is clear that the time for such a project is now.

In terms of the Grid's relationship with networking, it is important that there is tight collaboration between the networking infrastructure providers and the Grid middleware development community. Grid infrastructures desperately need L1, L2 and L3 end-to-end provisioning. Support for IPv6 is also an important factor in the future development of Grids.

In summary, the Italian Government fully supports the establishment at EU level of eInfrastructures in the context of the Grid and in agreement, and well integrated with, national initiatives. Only in this way will eInfrastructures help to strengthen the vision and delivery of ERA.

4.4.5 Walter Hoogland – Perspective from the Netherlands: a bottom-up approach

In the context of the Netherlands the presentation by Manuel Delfino had a clear resonance. The Dutch e-Science approach has been one of integrating a very high bandwidth infrastructure with the computer science community (interfacing to the physical network and the applications community), a number of typical e-Science applications, and extending this work to possible e-Business developments. The e-Sciencepark Amsterdam has become a focal point for the nationwide programme. It combines infrastructure (SURFnet and SARA) with the computer science environment and a potentially large number of advanced users.

The essential ingredients of this work include: the establishment of a framework in which Grid producers and consumers interact; the key issue of removing bandwidth constraints; finding a balance between technology push and applications pull; integrating networks and Grids; and producing a differentiated infrastructure capable of meeting the needs of high end applications and Internet users. Three large projects have been proposed for funding: the GigaPort, Virtual Laboratory for e-Science and GigaPort Next Generation Applications projects. Funding for these projects will be confirmed in September.

The Virtual Laboratory for e-Science will focus on creating an interactive problem-solving environment with a focus on methods and techniques for interactive High Performance Computing. It will also focus on adaptive information disclosure in the context of knowledge extraction, virtual reality based visualisation, collaborative information management and the integration of all of these components to form a virtual laboratory. Many typical e-Science applications will be supported. A key component of this framework will be the advanced networking research taking place in the Netherlands and focusing on the NetherLight Network, which is establishing an international lambda Grid.

In the long term a clear goal for the Grid should be its integration at the level of countries, disciplines, and academic and industrial users. This will require us to cope with different cultures and legal systems and in particular security issues. To establish a European Grid we may need to create a pan-European Grid service organisation but this will have to be different from such organisations in the networking context because with regard to the Grid there is clear value and ownership of local resources involved and control of resources is much more distributed. We must focus on the issues of portals, markets and Grid economics.

It is clear that it is in the interests of the Member and Associated States of the EU to support the creation of a national and pan-European eInfrastructure for e-Science. Arising from this there are a wide range of policy issues to address including: resource sharing, regulatory frameworks, brokering and security. In terms of the appropriate administrative and support schemes it seems sensible to focus these on honouring the rights of resource owners, identifying clearly who is responsible for what and observing existing contractual and trust relationships. In this regard the NRENs could play a major role. We should also consider in the same way as ERA, the establishment of a European Education Area, linking not just universities but also schools.

4.4.6 Aleksander Kuszniir – CEGC: a view from the local, regional and European perspective

The Central European Grid Consortium (CEGC) is composed of partners from six central European countries: Austria, Czech Republic, Hungary, Poland, Slovakia and Slovenia and was formed in 2002 in direct response to the establishment of the EGEE proposal. There are 13 partners currently involved in the consortium. CYFRONET is typical of these partners, with over 400 Gflops of installed computing power connected to the local metropolitan area network and also GEANT.

The CEGC partners have extensive experience in a wide range of Grid research and European cooperation and are involved in wide range of EU funded projects including: DataGrid, CrossGrid, GridLab, EUROGRID and GRIDSTART.

A typical CEGC project is the FloodGrid project, which is developing an interactive computing Grid aimed at the forecasting and management of flooding crises throughout Europe. This involves a cascade of time-critical simulations (meteorological, hydrological and hydraulic) and requires substantial computing power in “burst-type” activity cycles. The results of the project are currently being deployed for the Vah river basin in Slovakia a part of the CrossGrid project.

With regard to EGEE, CEGC has agreed to develop and maintain a joint Central European Grid Operations Centre, providing computing and storage resources as well as Grid access points and user services. A wide range of dissemination activities is also planned.

In FP6 the key for CEGC will be continued and expanded cooperation with partners from across Europe.

4.4.7 Marcel Kunze – Perspectives of Grids and e-Science in Germany

One of the key problems with the Grid in Germany has been that networking costs are very high and have to be born by the end user, leading to minimization of bandwidth consumption. Furthermore, the federal nature of Germany complicates interoperation and makes common national access to resources difficult. Both of these issues have resulted in a slow start to the national Grid initiative in Germany.

The German Grid Initiative held its kickoff workshop in February 2003. The 15 research centres of the Helmholtz Association and DFN-Verein initially drove the initiative that is open to all interested partners in academia and industry. Funding for the initiative is planned to become available towards the end of 2003 and is expect to be of similar scale to the other national initiatives in Europe.

A number of strategic considerations have become apparent. As the Grid infrastructure will integrate resources from science and industry, standardisation of interfaces and components and the implementation of a backbone network with non-trivial quality of service have emerged as key issues. The current DFN backbone generally operates at 2.5Gbps with a first 10Gbps line between Leipzig and Frankfurt (the German GEANT PoP). The need for the deployment of a network of competence and coordination centres has also become apparent and a network of Grid Support Centres will undertake this. At the same time the implementation of generic and prototypical Grid

applications in order to increase the user base is seen as crucial, and alongside this, the development of problem solving environments for e-Science and industry.

There are many technical challenges associated with building Grids and the approach has been taken of building a Grid “amongst friends” initially – with all sites deploying similar hardware and a long history of established trust. Of course, such a system will not scale and therefore thoughts are now turning to building a general public production eInfrastructure Grid. The key features of this include reliability and resilience, heterogeneity (in both hardware and programming models), and the provision of programming environments and tools for debugging Grid applications.

Building an international eInfrastructure will be even more difficult. A whole new range of issues will come to the fore including: logistical challenges – how do we maintain software, and enable accounting and billing across multiple administrative domains? Legal challenges – how do we deal with varying institutional policies and licensing models? Ideological challenges – how do we create a suitable framework to promote the development of stable persistent infrastructures? Can people make a profit from Grid services? Should Grids be centralised or distributed, free or charged, etc etc? Political challenges – including the integration of different cultures, the global, secure management of resources, decision making and enforcement and the provision of long-term stable funding to support the infrastructure. All of these issues must be answered over the next few years and this will not be easy.

4.4.8 Lennart Johnson – Grid deployment and support – the NGC, EGSC and SweGrid initiatives

The Nordic Grid Consortium was initiated a year ago and involves KTH, PDC and CSC who in this context act as service providers to the Nordic Grid community. Because each of the Nordic countries has a relatively small population, the aggregation of resources that is made possible is very important. One major issue that has arisen is how to share software – software licensing needs to catch up with the Grid concept.

The Consortium is basing its work on the need for a common security infrastructure and the need for a portal for job submission. Issues arising include the need for good security, appropriate resource sharing, new licensing models and data management middleware. The infrastructure is being built on top of regional networking initiatives, such as NORDUNET, which are already in place. While working in the area of Grids over the past 5 years, it has been interesting to see how Grid demonstrators have driven the need for network upgrades, from the then high bandwidth of 34Mbps to the 2.5-10Gbps networks we see today.

The NORDUGRID project is also an important Nordic activity. This project is based around the Nordic High Energy Physics community and has created strong links between this community and the DataGrid project. This has also led to the establishment of the Nordic DataGrid facility and the creation of the European Grid Support Centre in collaboration with PDC, CERN and CCLRC from the UK.

In Sweden the Government has been very slow to commit funding to the Grid. The SWEGRID project was funded at the end of last year and involved six Swedish Centres. It creates a Grid with nodes at each site consisting of around 100PCs and 20Tb of data storage. Clustered around this project are a series of Grid research projects focussing on resource management, distributed databases and security.

In the wider context, the real challenges for the Grid in Europe centre on the issues of authorisation, authentication and accounting – the “three As”. The concept of regional clusters, which had been mentioned several times earlier in the meeting, was also strongly supported.

4.5 Session 5: Application initiatives

4.5.1 Hans Falk Hoffmann – Grids and LHC: towards a first global Grid prototype

The challenges of Particle Physics over the next decade are key drivers for the development of the Grid. The Large Hadron Collider (LHC) will generate unprecedented amounts of data that will be analysed by researchers distributed throughout the world. Each of the four experiments located on the collider will generate multiple petabytes of data and in total nearly 500 institutes and over 5000 scientists will work together to analyse the data.

The Grid has been foreseen for some time. In 1992 Larry Smarr and Charlie Catlett wrote that “*Eventually, users will be unaware they are using any computer but the one on their desk, because it will have the capabilities to reach out across the Internet and obtain whatever computational resources are necessary*”. We are now beginning to see the realisation of that vision.

In the context of CERN, the EU DataGrid project has been key to the organisation’s involvement in the Grid. From the starting point of DataGrid a number of other projects have developed and formed around it. These include DataTag, CrossGrid and most recently the LHC Computing Grid project (LCG). The central goal of LCG is to prepare and deploy the necessary computing environment to enable the experiments to analyse the data coming from their detectors. In Phase 1 (2002-2005) the common applications, libraries and frameworks will be developed alongside a prototype environment that will involve the operation of a pilot-Grid service. In Phase 2 (2006-2008) the project will acquire build and operate the LHC computing service. The first prototype, LCG1, will be available from July 2003 and will be built from components of the DataGrid project, and VDT (which includes Globus and Condor). This prototype will provide low functionality but will be a real 24*7 service.

In the international context of LCG, FP5 and FP6 have been and are of the utmost importance. Likewise the UK e-Science programme has been very important in its commitment to Grids for Particle Physics and support of CERN – it has set the scale for other national initiatives to aspire to. LHC is a global resource and links to other national initiatives in the Nordic countries, Italy, Germany, France, Central Europe and South East Europe are also very important. Likewise links to, and support from, the US through the NSF Cyber Infrastructure and DoE Global Science Infrastructure programmes have been vital.

Recently, Paul Messina has proposed the concept of a Global Grid Middleware Institute. The mission of such an Institute will be to produce and maintain standards compliant and interoperable Grid middleware. It has been proposed that this Institute be a virtual organisation funded by the EU, European countries and several US Federal agencies, perhaps also involving the Asia Pacific region and industry. Its goal will be to ensure that Grid middleware becomes production strength and acquires sufficient functionality quickly enough to meet the needs of emerging Grid middleware user communities.

The EU set itself the goal at the Lisbon summit in 2000 of becoming “... *the most competitive and dynamic knowledge-based economy in the world ...*”. If Europe is serious about this goal then it will require considerably more funding and effort than is currently being expended. We need to up the ante.

In Geneva in December 2003 a world summit on the Information Society will be held. Currently this summit only involves Government and Business. This is a missed opportunity for the Science community. To redress the balance somewhat, a conference will be held at CERN on the 8th and 9th December, before the summit, to focus on the contribution of Science to the Information Society.

4.5.2 Fabrizio Gagliardi – Towards a common European market for computing and data management

The Grid vision is conceivable now because of the advanced state of computer and networking technology today. As a result, several software toolkits – Globus, Condor and Unicore – have been developed to achieve this vision and these are continuing to mature. A number of projects have demonstrated real early successes in various aspects of Grids. Europe has achieved a prominent position in this field in particular through the success of the European DataGrid project.

The DataGrid project was established in 2001 for three years with funding of €9.8 million and overall costs of approximately double this. A total of 21 partners are involved from research and academic institutes as well as industrial companies. Around 90% of the funding is allocated to the production of middleware for applications in the areas of High Energy Physics, Earth Observation and Genomic Exploration. The user community is continuing to grow and embrace new applications areas. Since last year the project has focussed on software quality (EDG 1.4.3 is the most stable release so far) and considerable increases in the number of geographical distribution of sites involved in the DataGrid testbed. The core testbed now consists of 12 sites from five countries who contribute a total of 1,075 CPUs and 15Tb of disk space. The testbed has made considerable use of the GEANT network, which has demonstrated excellent performance and is a major achievement for the EU.

A number of EU funded Grid projects have links to DataGrid and these include: CrossGrid, DataTag, Grace and GRIDSTART. DataGrid also has excellent links to a number of national initiatives such as the UK e-Science Programme, INFN-Grid and NorduGrid. However, as yet there are no real production quality Grids that can offer continuous, reliable Grid services to a range of scientific communities.

In the context of eInfrastructures, our vision must be to integrate current national, regional and thematic Grid efforts in order to create a seamless European Grid infrastructure. To exploit the Grid expertise that has been generated by EU supported projects and national Grid initiatives. To 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 geographical location. To provide a unique tool for collaborative, compute-intensive science (“e-Science”) in the European Research Area. Finally, to provide interoperability with other Grids throughout the world, including the US NSF Cyberinfrastructure – contributing to efforts to establish a worldwide Grid infrastructure.

Many obstacles lie in our way. In terms of technical challenges: current Grid middleware is often not interoperable – although a number of ongoing activities are seeking to remedy this situation; local site policies do not take Grids into account – security policies are not uniform, accounting procedures are specialised and not interoperable, and authentication and access policies do not allow for single sign-on. We need more uniform site policy standards.

In terms of political challenges: there are no business models to motivate industry – although all major IT companies are making strong statements and showing growing interest; the perception among traditional computer centres is that Grid technology could eventually undermine their market; the role of the commercial telecommunications operators and NRENs needs to be better understood; and standards for seamlessly connecting to the Grid, publishing information and bidding for resources are still emerging.

It is possible to make an analogy between the current state of Grid development and the emergence of the Internet. Early networks were largely incompatible. NFSNET (US) and JANET (UK) decided to provide network connectivity to their combined user bases. This large user base exposed security

holes and helped define common and acceptable use rules. From this the network we now know as the Internet emerged.

The EGEE project – Enabling Grids for e-Science in Europe – has as its central goal the creation of a production quality infrastructure built on top of the current and future EU research network infrastructure. It will build on the EU and Member States major investments in Grid technology and seek to build on international connections with the US and Asia Pacific regions. It will build on several pioneering projects – most notably DataGrid – and involve an established Grid development team of 60 people. Its overall approach will be to leverage current and planned national and regional Grid programmes (for example LCG) and work closely with relevant industrial Grid developers, NRENs and worldwide projects.

EGEE will have a major societal impact. For example: an international network of scientists will be able to model a new flood on the Danube in real time using meteorological and geological data from several centres around Europe; a team of engineering students will be able to run the latest 3D rendering programs from their laptops using the Grid; a geneticist at a conference, inspired by a talk she hears, will be able to launch a complex bio-molecular simulation from her mobile phone. Access to a production quality Grid will change the way science and much else is done in Europe.

There is also a political context to EGEE. The current Grid research and development projects will all complete within the next 18 months. The EGEE partners have already made major progress in aligning national and regional Grid research and development efforts in preparation for EGEE. Launching EGEE now will preserve the current strong momentum of the European Grid community and the enthusiasm of the hundreds of young European researchers already involved in EU Grid projects.

The key actions for Europe are: to establish a large production European Grid to support a Common European Market for computing and data management; to create an international board of senior stakeholders including representatives of resource providers, regulatory agencies and major user communities; to use this board to monitor and support the creation of an adequate regulatory framework for the Grid.

4.5.3 Mark Parsons – The Grid: challenging HPC infrastructure provision in Europe

Making large nationally funded High Performance Computing resources available on a European Grid infrastructure raises many issues for HPC providers. We must address these issues, which are largely of a political and policy nature now, if we intend to meet the vision of ERA and construct a real eInfrastructure for Europe.

EPCC, the supercomputing centre at The University of Edinburgh, has run major, nationally funded HPC systems for over a decade. From the early Transputer based Meiko Computing Surfaces, via the Cray T3D and T3E systems of the mid to late 90's to the Sun E6800 and E15000 systems and most recently the 1,280 processor IBM p690 based system (HPCx) which is currently Europe's largest HPC system for academic and research use.

Since 1993 access to these machines has been made available on an ad-hoc basis to over 400 EU visitors via our Training and Research in Advanced Computing Systems projects (TRACS). This project is funded by DG RESEARCH via the Access to Research Infrastructures action of the Improving Human Potential Programme. While the programme focuses predominately on access to large scale facilities, and in its latter incarnations has focussed exclusively on access, we have always provided support to our visitors to enable them to make best use of our systems and to learn transferable skills on their return to their home institutions. Throughout the project we have realised that the funding available is not enough to buy large-scale access to our HPC systems. These systems are not owned or paid for by the EU; UK funding has been used to purchase them to benefit UK scientific research. EPCC has made access available to these systems from its own time

allocation on the machines, which has generally been of the order of a few percent of the total machine capacity. In the context of HPCx this means over 50 users will get total access of around 0.5% of the machine. At the same time, visitors are only granted access to these systems for a limited period of time on return to their home institution – a key focus of our work has therefore been to ensure their codes are made or remain portable.

In general large HPC systems are bought to study problems infeasible on smaller systems, to focus on major scientific and engineering challenges. Crucially, they are not batch system replacements for workgroup servers. Access to these systems via the Grid poses many challenges including: authentication and authorisation issues; accounting software in the Grid context is in its infancy; there is currently no way to “trade” cycles; and security holes in rapidly changing software are a clear issue. The European Research Area complicates matters still further.

It is instructive to consider HPCx as an example. The UK taxpayer has purchased the HPCx system for UK scientists to use and thereby to benefit the UK economy. If, for instance, a German researcher wishes to buy time on HPCx because her problem won't fit on her IBM p690 system, they can offer cycles on their machine in exchange for access to HPCx. However, this is not a good deal for HPCx as the system was bought to solve large problems – UK users may not want their job migrated onto the German machine because it is too small for their needs or their work is confidential. Alternatively, money could be paid for access. Again, this is not a good deal for UK researchers as their access has been reduced to the machine and the amounts of money will be quite small – it would take a long time to save enough money for even one extra node for instance.

If the EU is serious about ERA they have several options. The EU could purchase an HPC system for researchers across Europe – helping to solve a common “chicken and egg” problem of only small numbers of non-UK users requesting access to our large systems because their availability is always so limited. Alternatively the EU could engage with national procurement projects. For example at the next UK procurement the EU could add 10% to the total funding therefore enabling a machine 10% larger to be bought. This approach would benefit everyone – it would give access to the machine to users across Europe without harming national investment and UK users would benefit by getting access to a 10% larger machine for 10% larger problems. Costs for this approach would be of the order €7-10 million. Finally, national governments could work together to purchase a system for Europe.

In summary, access to HPC resources over the Grid is a big challenge. The issues go well beyond technical matters and some of them go to the heart of what we mean by the European Research Area. The benefits to national user communities of ERA need to be much more clearly articulated by the EU. Many of these issues can only be discussed and resolved at an intergovernmental level facilitated by the EU.

5 Panel discussion and wrap-up

Following the final presentation, the meeting then took the form of an open discussion with a panel of guests which included: Spyros Konidakis, Mario Campolargo, Jean-Louis Picqué, Dany Vandromme, Vasilis Maglaris, Tony Hey, and Hans-Falk Hoffman. The discussion was wide ranging and sought to draw some conclusions and concrete actions from the meeting. The discussion is recorded below. The discussion started with a question from Enzo Valente.

5.1 Summary of discussion

E. Valente: A major concern for the networking community is the question of why many Grid presentations mention the need for a direct relationship with telecommunications operators and actual network links?

V. Maglaris: I interpret this question as meaning “why do people want their own fibre”? This is a clear challenge to GEANT to provide what users want. Without this engagement with users these calls will continue. But we must all remember that GEANT gives wonderful economies of scale.

S. Konidakis: What has been raised is the potential mismatch between the expectations of the two overlapping communities.

D. Vandromme: This may be directly related to the structure of FP6 integrated projects.

T. Hey: In the UK we have taken the approach of encouraging collaborative research between representatives of the Grid community interested in networking and UKERNA. This approach has worked very successfully to ensure both communities’ expectations are met.

F. Beltrano: An important issue is the coordination of all EU Grid activities funded within the context of FP6. In particular, what is the status between the EC Units F2 and F3? Is there now a combined vision?

S. Konidakis: The Commission is a single entity. Although it is arranged into an administrative structure, it is one organisation with one vision.

M. Campolargo: eInfrastructures sits within the Research Infrastructures Programme. There is no attempt within this programme to tackle all of the problems posed by Grids. What we are discussing at this meeting are policy issues relating to the components required to build Research Infrastructures of which middleware is clearly part. Many forums are clearly possible but our focus must be to discuss how we best create and operate Grid infrastructures.

J. Gruntorad: Going back to the original question, in the context of CESNET we have joined LambdaNet for the specific reason that a small number of users need very large amounts of bandwidth. They must make their own arrangements outside the GEANT framework in order to meet their specific needs. They do wish to work with GEANT to explore new methods of operation specific to Grids. The current approach has both positive and negative aspects. On one hand it offers excellent opportunities for their researchers but on the other it is a very expensive approach.

A. Kuszniir: A problem that is quite critical is how to attract groups from the wider society to the Grid – what are the applications drivers? How do we move from scientific to more general use? In this context, in Poland, the need for a Virtual Library concept built on top of Grid infrastructure is of great interest. Would the move to support such new user groups be viewed positively in terms of increased overall funding?

S. Konidakis: This is an important topic.

H. Hoffmann: The whole Grid infrastructure agenda only makes sense if we have users and applications that wish to make use of it. There are a number of immediately interesting areas that could be tackled – in particular the health and environment sectors.

S. Konidakis: We must try and build applications around Grids in novel areas.

V. Maglaris: In Greece the Grid is seen as a cost effective alternative to buying large monolithic supercomputers. As such, Grids are seen as being applicable in a wide range of application domains.

M. Delfino: In this context we must consider if our developments are revolutionary or evolutionary. One example of a revolutionary application is email. For this event we were all bombarded by large emails. If the Grid is to enable revolutionary applications then could it not be designed to enable us to quickly and simple create dynamic virtual organisations for instance purely for the lifetime of this meeting? If the Grid can enable the creation of such technology then virtual organisations may well come to be seen as the revolutionary “killer application” of the Grid.

M. Campolargo: It is clear that those present at the meeting are the people who really understand the challenges we face in the context of the policies required to build an eInfrastructure for Europe. We must plan how best to take our discussions forward. However, there are many issues that are outwith our control – for instance FP6 is already a reality and we cannot influence its broad thrust now. The need to attract user communities is a clear message coming from the meeting. Is there a way the Commission could play a role in this for instance? How could we stimulate it? A key issue is how we grow an eInfrastructure for Europe from the many national initiatives.

W. Hoogland: Agreed that there is already an obvious gap between advanced and new users. The key question is how do we attract more new users to the domain? We can clearly see the benefits of eInfrastructures to these domains but can the actual communities also see these benefits? There are clearly big cultural differences between the physical sciences and other sciences. We need to win the hearts and minds of individual scientists across a wide spectrum of domains. The possibilities for eInfrastructures are large, the EU must assist with regional centres to help them achieve this goal of building larger and broader user communities.

V. Maglaris: In order to create a real Grid we need to establish a web of trust. This is vitally important or we will not see the major benefits of Grids. We run the risk of merely establishing multiple standalone clusters. We need to create trust in the same way that the networking community has.

H. Hoffmann: We need to generate more applications to make use of the infrastructure and a middleware that can quickly react to their needs. We must focus Grid development on what people want from it. This is a good argument for the need for an Open Grid Infrastructure Institute.

Norwegian rep.: In Norway we are witnessing a sharp division between basic research and applied research and this is a clear problem. The applied research community view Grid building as an issue for the basic research community and vice versa.

S. Konidakis: We need greater promotion of the benefits that eInfrastructures will bring to many application domains and ensure future developments are driven by applications pull rather than technology push.

M. Mazzucato: Responding directly to Mario Campolargo’s points: there is a clear need for a Policy Forum. This should be driven by the research part of the community but at present we are not seeing a strong push in this direction. We need to stimulate the creation of critical mass. At present a clear issue is that potential applications providers do not contact us – we contact them. A real critical mass behind the Grid would improve this situation but we need to ask the question – how do we achieve this?

S. Konidakis: Invited other comments on the idea of a Policy Forum.

M. Parsons: A Policy Forum is clearly a good idea but we must be careful to construct it so that it can achieve its goals. We should make sure that all of the regions involved in eInfrastructures are properly represented at a high level. This may involve inviting people from a Governmental context. We should aim to hold 2-3 meetings of the Policy Forum per year and we should ensure that a small number of people, funded to work within the Forum, produce White Papers in the intervening periods which can then be debated, amended and endorsed by the high level attendees at the Policy Forum meetings. Without these people, the members of the Forum will not have enough time to produce the required documentation and the goals will not be achieved.

W. Hoogland: A clear concern with regard to the Policy Forum proposal is that we should be very careful about not inviting people who are too high level. Nothing will be achieved if we pick the wrong group of people.

M. Delfino: Following on from the previous statements. We also need to be very careful that the subject matter is carefully set. Although high-level people may have decision-making authority, they may not have the correct level of knowledge to adequately debate the difficult issues we have been discussing today. What needs to be discussed at a high level is that EU science is under-digitalised and because of that collaboration is not sufficiently enabled. Posed in this way, high-level policy can be formed and then those lower down the chain can build on this to form clear specific policies to achieve this goal.

M. Campolargo: Pointed out that some high-level committees already existed.

J. L. Picqué: In the context of DG RESEARCH's multi-disciplinary bottom-up programme on Research Infrastructures, researchers are applying to develop "Grid-like" digital libraries. This is still very new but it is clear that some domains – in particular Particle Physics, Astrophysics and Bioinformatics – are clearly leading the way. Focusing on involvement in this process from a high-level, the European Strategy Forum on Research Infrastructures, or ESFRI, has recently been created to discuss policy-making on Research Infrastructures. This Forum is now planning to start a Working Group, including national experts, to discuss networks and Grid issues. The Working Group will report to ESFRI. In the context of this meeting it would be very interesting to understand how ESFRI can properly discuss a coherent approach in Europe to the challenges we have been discussing at this meeting. At a slightly lower level, we should perhaps consider how best to involve the people represented at this meeting in Working or Advisory Groups, with representatives drawn from each country, in order to establish coherent policies. It is useful to consider in this discussion how there are two levels – the political and the scientific.

M. Campolargo: Agreed with Jean-Louis Picqué's analysis. To add to his comments, it is clear that every activity represented by the Commission has different needs. For instance an Expert Group has recently been launched to establish a technology roadmap for Grid research in Europe. In the Research Infrastructures community in Europe we do not have a lack of expertise – what we lack is people taking responsibility in Europe for engaging in a detailed dialogue with the demanding user communities. In essence we are missing an eInfrastructures Reflection Group that could prepare White Papers and meet with and brief Governmental representatives. There is a clear gap here.

V. Maglaris: We need to be very careful how we discuss the establishment of new forums. We don't want to make the mistake of replicating work going on elsewhere. For instance we wouldn't want to replicate the work of the high-level "thinking" groups within the EC. The objectives of this meeting included the definition of a policy initiative for Europe and the creation of a high level committee not just for Grids but also for networking – the concept of eInfrastructures. The whole concept of eInfrastructures will need to be carefully and clearly

explained to the Council of Ministers. It was proposed that a representative group of this meeting meet to advise the people who participate at a high level within the Commission. We need to focus this approach very carefully – such a committee needs a clear remit, agenda and role.

S. Konidaris: We need to carefully decide, as a result of this meeting, what level of committee or committees are required and what exactly their remit should be.

M. Delfino: In summary – we must focus on end-to-end delivery of eInfrastructures.

H. Hoffmann: In terms of applications we should consider which have a clear European dimension and encourage the European Commission to promote such applications at the European level. In this context Healthcare comes to mind.

S. Konidaris: We would need to think very carefully about this as the Commission must always be seen to be even handed.

J. Sanchez: We should emphasize the need for continuity of the decisions of this workshop and we also must not forget the word “Beyond” in the title of this meeting – we must not forget industry and moving forward in this direction which may also directly effect the policies we must put in place. European industry should be involved and be able to acquire the knowledge generated from the beginning.

As this was the final point made during the discussion the Chairman and the Rapporteur summed up the main recommendations from the meeting which are presented in Section 6. Following the summary of these recommendations a number of final points were made which are summarised here.

M. Campolargo: As we are nearing the end of the Greek Presidency of the EU, a troika of the Greek presidency and the two following ones (Italy and Ireland) should discuss further what needs to be achieved with regard to moving this debate forward. We must carefully consider our next moves. One important factor in the success of this meeting has been the blessing of the Greek Presidency for it to be held. A second meeting would be very useful and it is hoped to hold it in the context of the Italian Presidency.

S. Konidaris: We need to better articulate what we mean by eInfrastructures if we are to achieve our objectives and clearly explain our goals at a high level within the EU. Perhaps a short document or note is needed to express what we have discussed today. We must articulate our vision at all levels.

6 Recommendations

The meeting approved the following recommendations:

- The European Research Area should clearly be seen to embrace Innovation – articulated in the context of this meeting through the name European Research and Innovation Area (ERIA).
- The strong level of interest in the meeting indicates how eInfrastructures are vital for the attainment of the vision of eEurope and ERA.
- It is clear that many countries are joining together into *regions* and this was presented as a powerful tool for cooperation. An EU-wide infrastructure could grow from these regions.
- eInfrastructures will only succeed if we solve end-to-end issues at the technical, infrastructural, methodological and social/human levels.
- GEANT is a major achievement and may show the way forward in terms of building production Grids and a real eInfrastructure throughout Europe.
- Solving the challenges of *authorisation*, *authentication* and *accounting* are key challenges for all Grid projects – this is a major hurdle in the context of building an eInfrastructure for Europe.
- The trust model has to be developed further in order to share not just bandwidth but also computing resources. Grids must take the lead in helping with this process.
- The next steps for the Grid must be to move to *reliable*, *resilient*, and *robust* production quality middleware.
- We should continue to focus on Open Standards and avoid any vendor lock-in.
- The idea of an Open Middleware Infrastructure Institute for Europe was broadly supported – the rationale behind this being to create the next generation of production quality software from the developments that have taken place to date.
- Key to the general uptake of Grids and the creation of a real eInfrastructure for Europe will be the transition from e-Science → e-Business → e-Society.
- We must identify the next generation of applications – the so called “killer apps” – and improve our promotion of the benefits that eInfrastructures will bring to their user communities.
- To build eInfrastructures we need to focus on middleware interoperability and the accompanying policy decisions required to make our software and operating paradigms interoperable in a global context.
- Policy issues – particularly in a *local* context need to be addressed. Only by addressing the intricacies of local policy issues will be able to make local resources available in Grids.
- We will build eInfrastructures by focusing on policy issues related to resource sharing in the context of the European Research Area. Such discussions must take place at an intergovernmental level.
- The overall recommendation from this meeting is that an eInfrastructures Reflection Group, built from National Programme representatives, should be established and perhaps advise the Governmental representatives who sit in existing committees.
- The eInfrastructures Reflection Group should consider and communicate clear messages on eInfrastructure Policy issues to both the European Commission and existing eInfrastructure projects on policy matters.

- A troika of the current presidency of the EU (Greek) and the two following ones (Italy and Ireland) should discuss further what needs to be achieved with regard to moving this debate forward.

7 Closing

As this was the end of the meeting, the Chairman, Vasilis Maglaris, thanked all of the attendees for their time and effort and also Spyros Konidaris and Mario Campolargo of the European Commission for their input, support and Chairmanship. He also thanked DG RESEARCH and DG INFSO for their support and vision. A final vote of thanks was given by Mario Campolargo to Kyriakos Baxevanidis, Jorge-A. Sanchez-Papaspiliou and all of the staff at GRNET for their excellent organisation of the event. He invited the Greek Presidency to consider turning the meeting report and presentations into a short brochure, which could act as a milestone in the development of this important area.

Dr Mark Parsons
EPCC, 20th June 2003