e-Infrastructure Reflection Group

White Paper

Version 1.2, 30/6/2006

This document contains the White Paper composed during the Austrian Presidency of the European Union and the Austrian Chairmanship of the e-IRG. It provides an updated view of the status quo, assembled by e-Infrastructure experts upon appointment by the e-IRG, which will serve as a basis for future decisions and possible actions. The plan contains contributions on: Authentication And Authorisation Policies, Incentives for Resource Provision in Scientific Grids, Usage Policies, Training and Education, Grid Economy – Allocation and Accounting, Middleware Interoperability and Repositories, User Support, Supercomputing and Grid Harmonisation Policies, Responsive Grids – Short Jobs and Policy Issues, and Legal Issues. The material was elaborated and edited following the discussions that took place in the e-IRG workshop in Linz and after integrating comments from the e-IRG meeting in Vienna.

AUSTRIAN GRID
bm:bwk GZ: 4003/2-VI/4c/2004
http://www.austriangrid.at
contact@austriangrid.at

e-Infrastructure Reflection Group
http://www.e-irg.org
secretariat@e-irg.org

CONTENT

1. PREFACE

Since its original initiation in 2003 during the Greek EU presidency, the e-Infrastructure Reflection Group (e-IRG) has been able to demonstrate its importance as a contributor and facilitator in the area of e-Infrastructures and related topics. At this point in time, we are making the transition from the 7th e-IRG leadership during the Austrian EU Presidency to the 8th leadership during the Finnish EU Presidency, which provides me with an opportunity to reflect on our achievements.

The developments in e-Infrastructures over the last three years have been astounding for members of our community, observers and "customers" alike. Thanks to all the e-Infrastructure contributors around the globe and in Europe in particular, e-Infrastructures such as the research networks, the grid infrastructures, the high performance computing infrastructures and the digital libraries have experienced a massive push in many different aspects. This clearly drives the products to further and further maturity, thereby promoting the deployment and adoption of e-Infrastructures for the benefit of science and research, and society.

At the same time, the level of maturity of e-IRG itself has been raised. The workload required to perform the mission of e-IRG, namely "to support on the political, advisory and monitoring level, the creation of a policy and administrative framework for the easy and cost-effective shared use of electronic resources in Europe (focusing on Grid-computing, data storage, and networking resources) across technological, administrative and national domains" has required substantial efforts from the e-IRG Delegates and Observers, the e-IRG leadership and the support personnel. The e-IRG has dynamically adapted itself to these changing needs, and further optimisation and streamlining will occur during the next couple of months. With the agreement on a more permanent leadership model, e-IRG is looking forward to a modus operandi which guarantees the best possible support for the needs of the e-Infrastructure community in the interests of its customers.

As another concrete result of the Austrian e-IRG chairmanship, efforts have been made to identify necessary actions for enabling and guaranteeing the sustainability of e-Infrastructures. With the setup of a focused task force, and the discussion and finalisation of the paper on "Sustainable e-Infrastructures" (SeI), the e-IRG has identified a set of five recommendations to make e-Infrastructures more sustainable. The resulting paper was well received by the European Commission and is now being communicated to the attention of the individual national funding agencies and policy makers, to provide guidance in future investments in e-Infrastructure across Europe.

This version of the e-IRG White Paper, the e-IRG Linz White Paper put together during the Austrian chairmanship, represents another landmark in the work of the e-IRG Delegates, the e-IRG Observers, and the nominated experts working on the contents of the paper. The White Paper provides an updated view of the status quo, assembled by e-Infrastructure experts upon appointment by the e-IRG. Therefore, it provides the basis for future discussions and actions in our domain. The recommendations in this paper identify emerging issues and policies that will be studied and evaluated by the e-IRG during the upcoming months. Wherever appropriate, key aspects may be formally developed into recommendations, formally endorsed by the eIRG. For this reason, the e-IRG White Paper represents a valuable contribution as a status report and discussion paper, while at the same time identifying tasks for the future work in this area.

Clearly, there are many open and intriguing questions and challenges ahead, and the e-IRG is looking forward to an inspiring process of monitoring, evaluating and encouraging the development of the e-Infrastructures. We should not forget the importance of working together to achieve this goal, and I am convinced that e-IRG will be able to continue its successful cooperation in the future.

Personally, I am extremely grateful for the active involvement of all the e-IRG Delegates and the input and guidance from our Observers, a cooperation which has continually driven the vision and the work of the e-IRG. The results achieved during the Austrian presidency, including this particular document, are the fruit of the active contributions of many individuals. I want to thank all who contributed to the work of the e-IRG during the last couple of months, either through active participation in the e-IRG Open Workshop or the 6th and 7th Delegates Meeting, or more directly through providing input, ideas and comments via various means of communications.

My personal contribution to e-IRG was only possible due to the support provided by the Austrian Grid Initiative and the Austrian Federal Ministry of Education, Science and Culture (BMBWK), my fellow Austrian e-IRG Delegate, and the colleagues from the e-IRG Troika. Last but certainly not least, I would like to thank the e-IRG White Paper Chief Editor, Fotis Karayannis, and the members of the support team, Christian Glasner, Joanne Lawson, Matti Hekkurinen, Martin Polak, and Panagiotis Louridas for their hard work.

On behalf of the e-IRG,

Dieter Kranzlmüller

(Austrian e-IRG Chairman)

2. INTRODUCTION

2.1. PURPOSE OF THE DOCUMENT

This document contains the White Paper that produced during the Austrian Presidency of the European Union and the Austrian Chairship of the e-IRG.

It was decided that the White Paper would be produced in two steps. First, an initial version would be produced, which would the subject of discussions during the e-IRG Linz workshop. This version was called “White Paper Concept”. Based on the outcome of the discussions, the White Paper Concept would be amended to produce the final version of the White Paper.

This White Paper version was produced on the e-IRG wiki, which is hosted at http://wiki.e-irg.org. The e-IRG decided on the sections that would comprise the White Paper Concept, and an outline was provided to the White Paper contributors with the help of the e-IRGSP (the outline is available in Annex 4.2). The contributors were then guided by the e-IRGSP on how to edit their material.

This documents contains the editors' contributions, as these were shaped by the discussion that followed the White Paper Concept.

It must be noted that the e-IRGSP provided logistical help and guidance during the preparation of the White Paper Concept, it was and is responsible for the wiki on which it is edited, but the contents are the fruits of the contributors' efforts. The contributors are listed in Annex 4.1.

2.2. DOCUMENT AMENDMENT PROCEDURE

The White Paper concept is the result of close cooperation between the e-IRG and e-IRGSP; changes to this document should be approved by the e-IRG.

3. EXECUTIVE SUMMARY

This document contains the White Paper composed during the Austrian Presidency of the European Union and the Austrian Chairship of the e-IRG. It was decided that the White Paper would be produced in two steps. First, an initial version would be produced a so called "White Paper Concept", which would be the subject of discussions during the e-IRG Linz workshop. Based on the outcome of the discussions, this initial version would be amended to produce the final version of the White Paper.

This White Paper version was produced on the e-IRG wiki, which is hosted at http://wiki.e-irg.org. The e-IRG decided on the sections that would comprise the White Paper Concept, and an outline was provided to the White Paper contributors with the help of the e-IRGSP (the outline is available in Annex 4.2). The contributors were then guided by the e-IRGSP on how to edit their material. This document contains the editors' contributions, as these were shaped by the discussion that followed the White Paper Concept and the comments received by the e-IRG members during the Vienna e-IRG meeting.

The issues discussed are:

- Authentication And Authorisation Policies

  This topic has been dealt with in the past White Papers providing already a series of recommendations. In this section efforts are concentrated on the convergence strategies of different AA technologies and frameworks. Forums such as the IGTF and the TERENA Task Forces steadily contribute in this direction and thus they should be encouraged. A new topic on traditional security aspects focusing on incident handling and response coordination has been produced by the editors as a theme for the next workshop. Note that the editors propose that the AA section recommendations are endorsed by the e-IRG.

- Research Networking Policy Challenges for e-Infrastructures

  This section highlights the role of the RN as the basic building block of the e-Infrastructures providing stable services to its different categories of users. RN policy should be based on well defined, mutually accepted principles and practice of sharing responsibilities and duties wrt. governance and control between the RN service providers and the RN users. Sustainability issues are then tackled proposing the investigation of continuous flexible funding schemes, while the role of foresight activities such as EARNEST are presented. At the end the wide range of requirements on network features that have to be taken into consideration by the RN policy strategies are given, while the role of education and training is then underlined. The idea of carefully extending the RN network to schools is also proposed.

- Incentives for Resource Provision in Scientific Grids

  This section appears in the White Paper series for the first time, having a research flavour, and thus should be handled appropriately. The main outcome of this section is that allowing free access to grid resources does not make sense (at least from an economics point of view), and thus governments and NGIs should consider putting a minimum threshold in resource contribution of a participating oraganization in order to avoid free-riding. This also seems to be supported by the observation that very small clusters, due to the administrative and technical overhead involved, may be counterproductive.

- Usage Policies

  This section topic has been examined in many of the previous White Papers providing the state of the art in AUPs and reflecting mainly the joint work of EGEE/LCG/OSG on a Grid AUP that the user should accept by registering with a VO. The latest version of such an AUP is included in the section, while a series of open issues have been highlighted related to the use of AUPs by other communities such as the Supercomputing and Digital Libraries ones, as well as commercial entities and SMEs. The use of an AUP by short-lived VOs was also examined and it was recommended to first study related use cases.

- Training and Education

  Training and Education is an area that has been proposed to be included in the White papers during the UK presidency and now appears for the first time. However, the role of training and education is universally accepted and has been embraced by different forums such as GGF, EC projects (EGEE, ICEAGE) and related "concertation" activities. In the section a series of issues are analysed such as characteristics of courses, mechanisms for developing standards and best practises, and a series of policies for the related trainers, the trainees, the material and the training infrastructures. The section concludes by recommending increased investments in training and education, as well as developing the corresponding policies.

- Grid Economy - Allocation and Accounting

  This section has been examined in previous White Paper versions, but this is the first time that goes deeper into accounting and economic modelling, besides the state of the art of research efforts. It examines different allocation processes of resources and concludes that there is no need for a single allocation service -a fact which is ignored by the current research efforts- rather than scalable and interoperable solutions. Corresponding economy models need to be developed. Accounting/metering is gradually becoming available in Grids and thus need to be standardised. Appropriate planning is also recommended for the future more sustainable Grid projects.

- Middleware Interoperability and Repositories

  This section first sets the issues that need to be tackled at policy level, namely the creation of software repositories, the packaging of software into appropriate distributions, the use of Service Oriented Architectures, and the promotion of quality assurance and standards compliance. Important efforts working on the above areas are summarised and the proposed approach to achieve the above mentioned goals as part of EU-OMII is given. EU-OMII will build on top of existing efforts such as EGEE, OMII-UK, ETICS, Globus and NMI, providing an open repository reporting on the software quality and interoperability.

- User Support

  The user support section provides the updates from the Luxembourg version with corresponding reviews and comments. It highlights the strong connection needed between training and user support in a follow-up role. Regarding the application support it concludes that this is linked with the size of the corresponding VO. Then the role of NGIs in addressing user support is put forward keeping the federated model, along with with the need for central coordination. Interfaces between Grid and networking user support structures are also covered, and the need to pursue the user support area in GGF is stressed. The editors note that their views are still based on the experience gained by the EGEE project.

- Supercomputing and Grid Harmonisation Policies

  This section proposes that Europe should establish a two-tiered e-Infrastructure of supercomputing centres: the existing national centres will form the Tier-2 centres. Tier-1, the tip of the new resource pyramid, is constituted by a few sites that will host shared European petascale supercomputers. The two tiers should be strongly integrated together into a global infrastructure "à la DEISA" with a unique operational and service providing model. Europe needs a very restricted number of high end petascale systems in the next few years to substantially enhance capability computing and to match the planned evolution in the US and in Japan. Supercomputing Grid infrastructures "à la DEISA" are not a substitute for high end petascale supercomputers, as they are complementary. Two important issues having policy implications are presented, namely the high performance access to remote data from among the different tier centers, and the delicate issue of funding, and the operational model of Tier

1s. The section finally proposes the reinforcement of research into new programming models, new algorithms and new methods that will enable the above-mentioned set of services to be established.

- Responsive Grids - Short Jobs and Policy Issues

  This section appears in the White Paper series for the first time, and as such it has a research flavour and should be treated accordingly. The section provides example scenarios where responsive grids can be beneficial. Then experience gained in past projects such as CrossGrid is given, while current efforts such as GridCC and int.eu.grid are mentioned. The section examines whether it is possible to define a roadmap for the integration of responsive grids middleware into the production one, which is a delicate issue. It proposes that this could be tackled by the e-IRG and/or ESFRI roadmap documents. Finally fruitful discussions should be triggered between the different stakeholders.

- Legal Issues

  This section follows up the recommendations of the Luxembourg White Paper and proposes the organisation of a corresponding workshop at the end of 2006 - beginning 2007. In this workshop legal experts in the field of information technology and law will analyse the specific features of e-infrastructures. Input to this workshop will be given by experts in the field of the technical and organisational aspects of e-infrastructures. It also proposes that an inventory with legal problems is made. It recommends that practical solutions for the legal problems currently being experienced be sought. These could range from (European) regulatory initiatives to standard-setting, developing certification schemes, codes of conduct, model contracts and other forms of self-regulation. In all cases, the issues should be addressed at the European level, and in some cases maybe even at a global level (treaties, standardisation initiatives, declarations etc.)

- Integrated Data Management

  A task force to tackle the above area has been created during the UK presidency. A first draft paper has been presented during the UK e-IRG meeting in London and the task force continued its work during the Austrian presidency producing an updated version of the paper. During the Linz workshop a report was presented to the audience. However, it hasn't been possible to adapt the Task Force report to the template of the White paper and include it in this section for this version. This will be done later this year, during the Finnish presidency, and thus this sections acts as a place holder for the future work.

- Terminology

  A first attempt to harmonise some of the basic terms in e-Infrastructure is given in this section. However it is more than certain that further work would be needed during the next presidency.

4. RECOMMENDATIONS

The list of recommendations provided in this section are taken from the corresponding sections of the white paper. It is felt that the list in its current form cannot be endorsed by the e-IRG, mainly because of differences in the level of maturity, the length and the style of writing. The issue was discussed in the e-IRG bodies and it was decided that the recommendations need to be worked further during the Finnish presidency. It has to be noted that some of the editors proposed that their section recommendations be endorsed during the Austrian presidency, but this was not feasible. The current plan is to endorse the recommendations after being further worked out during the Finnish workshop, 4-5 October, where a separate e-IRG meeting is envisaged http://www.e-irg.org/meetings/2006-FI/workshop.html.

4.1. AUTHENTICATION AND AUTHORISATION POLICIES

- Trans-disciplinary (Grid projects, NRENs, other user communities) and trans-continental forums that move towards the establishment of a global, seamless AA infrastructure for e-science applications should be encouraged. The e-IRG wishes to acknowledge the efforts made in this direction by the IGTF and the open information exchange point provided by TERENA task forces.
- As stated in the introduction, and as an additional recommendation, the e-IRG would like to introduce a new topic on security incident handling and response coordination in the next e-IRG cycle.

4.2. RESEARCH NETWORKING POLICY CHALLENGES FOR E-INFRASTRUCTURES

- A. Policy requirements
- A1. Research Networking policy should be based on the postulate that the Network is a dominant major element of the e-Infrastructure, supposed to provide perfect service both for less demanding everyday users and most demanding Grid users.
- A2. RN policy should be based on well defined, mutually accepted principles and practice of sharing responsibilities and duties wrt. governance and control between the RN service providers and the RN users. The proven RN governance model is available also for RN usage based on realistic user demands and on a permanent service provision model.
- B. Sustainability and development issues
- B1. Continuous EC financing rather than running finite lifetime EU FP projects in RN would be preferred, together with sustainable operation of NRENs and step by step easing the digital divide. Elevated EC funds (>0,5 % of the entire FP7 budget) in 2007-2013 and changing the 50 % rule are most desirable.
- B2. RN development policy should be matched to the emergence of promising/proven new solutions rather than to the periodicity stemming from the duration of the related EC projects. Smooth transitions and keeping leading edge position in global scale are major requirements. Planning for such developments needs also user inputs, eg. through the EARNEST activities.
- C. Requirements on network features and on user information
- C1. RN policy should wisely take into consideration a number of different aspects: technology options, network architectures, management/interoperability, AAAI options, user communities, areas of activities, business/economy, global connectivity, as well as efficient allocation/integration of e-Infrastructure resources, common language of networkers/users, common practice of network domains, Grid aspects, QoS, AUP, etc.
- C2. Informing, training, and educating the users should provide necessary knowledge about the RN and also about the Grid infrastructure. Easily accessible data/knowledge-bases, d-

libraries, e-repositories, e-archives should serve as resources of the requested information. Education and training, in general, should carefully consider extending the RN to secondary/ primary schools, by taking into account the differences between academic and teenager users.

4.3. INCENTIVES FOR RESOURCE PROVISION IN SCIENTIFIC GRIDS

- Allowing free access to grid resources does not make sense (at least from an economics point of view), and thus governments and NGIs should consider putting a minimum threshold in resource contribution of a participating oraganization in order to avoid free-riding. This is also seems to be supported by the observation that very small clusters, due to the administrative and technical overhead involved, may be counter-productive.

4.4. USAGE POLICIES

- The e-IRG encourages wider use of the proposed Grid AUP to promote interoperability and feedback on cases where the proposed AUP does not work saying what needs to change.
- The e-IRG sees a need to investigate how anonymity and/or pseudonymity can be introduced in the Grid environment to better protect the privacy of the Grid user (important in the medical field).
- The six open issues described above all need to be addressed. In particular, the uses of the e-infrastructure which do not fit the Grid VO model are likely to need a different AUP.

4.5. EDUCATION AND TRAINING

The e-IRG recommends that:

- investments in education and training should be made in order that the full potential of e-infrastructures may be realised for EU citizens;
- significant investment in education should be directed towards improving understanding as to how e-infrastructure may be exploited effectively;
- investment in training and education should be directed towards improving the ERA's capacity to develop, deploy and manage e-infrastructure and the applications that use e-Infrastructure;
- policies and standards for the coordination, collaborative support, interchange and mutual recognition of e-infrastructure should be developed across the ERA.

4.6. GRID ECONOMY - ALLOCATION AND ACCOUNTING

- The e-IRG considers that scalable resource allocation mechanism(s) must be included in the design phase of any Sustainable European Grid Infrastructure;
- The e-IRG supports the development of a versatile Grid economy model supporting those allocation mechanism(s) for a large number of small VOs, which should influence the design and implementation of any scalable and Sustainable European Grid Infrastructure;
- An accounting policy and standard that enables interoperability must be developed immediately, but be based on the vision of a full Grid Economy.

This work should be based on an information exchange and collaboration between the e-IRG at the policy level, the Global Grid Forum (GGF) for standardisation efforts, and accounting system developers to make sure that implementation considerations are taken into account.

4.7. MIDDLEWARE INTEROPERABILITY AND REPOSITORIES

- The e-IRG strongly supports the continued harmonisation of software services and tools contained in Grid middleware distributions through their engagement in standardization activities driven by practical implementation experience. Self-certification of standards compliance is not deemed sufficient, and the e-IRG considers the provision of an independent documented process for standards compliance and quality assurance to be vital. The e-IRG considers that the provision of a repository to include software components (with documented quality assessments) from other EU projects is an important area which should be discussed at future meetings.. It would provide a showcase for the EU community, allowing the the software outputs from various projects to be made available from one location.

The currently funded work within OMII-Europe and ETICS is considered to be crucial in this respect.

4.8. USER SUPPORT

We now review and comment on the relevant recommendations from the Luxembourg White Paper (which are in italics)

Infrastructure planning should include the provision of resources for the co-ordination and delivery of an educational programme covering all aspects of Grid usage.

- A strong connection between education, training and user support is needed!
- New users need user support after training.
- The coordination of information for Grid users is a challenge to be met.
- We must establish generic connections to middleware providers in order to interact with the suppliers of components.

We recommend that the infrastructure planning includes the provision of resources for the support of application areas, both in the form of support for the first application migration, and with dedicated ongoing support for application areas according to their scope and complexity.

- The resources and input you can expect from a VO in this area depends on the size of the VO.

We recommend that the developments of day-to-day support systems for the Grid be fully supported.

- The federated model of EGEE should be pursued.
- The National Grid Initiatives (NGIs) should contribute to a European Grid User support system.

It is important that user support system developments are fully documented and discussed in the Global Grid Forum (GGF) and with other major Grid projects. This applies also to the development of information systems for Grid users.

- We must pursue this area in GGF to ensure sharing of experience and technology.

A policy for networking support for Grid infrastructures should be agreed and made mandatory for all entities contributing to the operational environment. The implementation of the networking support should be interfaced to the overall user support infrastructure.

- Connections are already established, but more discussion is needed .
- Is there a centralized user support in networking?
- Connection to supercomputing centres/projects?

The e-IRG has reviewed the recommendations for Grid user support in the key areas of user education, the provision of easily accessible user information, support for applications and the day-to-day running of the Grid and networking infrastructure. The e-IRG would like to stress that the planning for such support of the infrastructure should allow for the continuity of support in national structures beyond the end of major Grid projects.

- The NGIs should address user support in order to ensure sustainability of e-infrastructures.
- NGIs should be coordinated on European level.

We recommend that the developments of day-to-day support systems for the Grid be fully supported.

- National Research and Education Networks (NRENs) should become more involved in the process to provide day-to-day Grid user support. This is in progess: GGUS and the NRENs are working on a collaboration between network and Grid support infrastructure through an interface between GGUS and the ENOC (EGEE Network Operations Centre).

4.9. SUPERCOMPUTING AND GRID HARMONISATION POLICIES

- The e-IRG recommends that the existing Grid supercomputing infrastructure – based on the strong integration of national services – be completed with a restricted number of leading petascale systems, oriented to a 'two tier' architecture for the global European HPC services
- The e-IRG encourages the reinforcement of research into new programming models, new algorithms and new methods that will enable the above-mentioned set of services to be established.

4.10. RESPONSIVE GRIDS - SHORT JOBS AND POLICY ISSUES

- The e-IRG recommends that awareness about the benefits of responsive Grids be raised, and that the above-mentioned issue of lack of immediate return on investment be alleviated by the creation of adequate incentives.
- The e-IRG believes that coordination between the different EU projects addressing the issue of Grid responsiveness as a research and development task should be developed. The goal is to promote software interoperability, and factorization of developments.
- A dialogue should be initiated between the Grid providers endorsing the objective of responsive Grids and the EU projects and actions in the areas of ubiquitous computing and ambient intelligence.
- At the technical level, the virtualisation frameworks and tools and their impact on responsiveness should be surveyed. Because industry investment in this area is already considerable, close interaction with industry software providers is required.

4.11. LEGAL ISSUES IN E-INFRASTRUCTURES

- The e-IRG proposes the organisation of a workshop aimed at identifying the legal issues specific to e-infrastructures. In this workshop legal experts in the field of information technology and law will analyse the specific features of e-infrastructures. Input to this workshop will be given by experts in the field of the technical and organisational aspects of e-infrastructures. The outcome of the workshop should support the decision-making process regarding the way that development of the legal component of e-infrastructures should progress.
- On the basis of the outcome of the Amsterdam workshop, the e-IRG recommends that an inventory is made of the legal problems which the scientific community currently feel impede the practical use of e-infrastructures.
- The e-IRG recommends that practical solutions for the legal problems currently being experienced be sought. These could range from (European) regulatory initiatives to standard-setting, developing certification schemes, codes of conduct, model contracts and other forms of self-regulation. In all cases, the issues should be addressed at the European level, and in some cases maybe even at a global level (treaties, standardisation initiatives, declarations etc.).
- The e-IRG believes it is time to start a more fundamental analysis of the legal issues involved in sharing computing resources and expanding the scope of this practice beyond the scientific community, both on the demand side (users) and the supply side (those making resources available).

- The e-IRG notes that legal issues should be an integral part of a multidisciplinary approach in the further development of e-infrastructures. Incorporation of legal aspects in the debate will be beneficial to the further development of e-infrastructures.

4.12. INTEGRATED DATA MANAGEMENT

This section acts as a place-holder for future recommendations that will be prepared by the Integrated Data Management Task Force during the Finnish Presidency.

5. AUTHENTICATION AND AUTHORISATION POLICIES

5.1. POLICY AREA

Authentication and Authorisation (AA) systems are a critical resource in current e-science infrastructures, and great progress has already been made in formulating policies and providing sound and secure methods for establishing and exchanging identities and rights.

5.2. GOAL

There have been significant steps in the direction of achieving the promise of global, seamless authentication and authorisation mechanisms, although there are still important challenges faced by architects and practitioners. This section analyses the current state of play, detailing achievements, obstacles and possible solutions.

It is also worth noting that during the process of preparing this white paper, the issue of the multiple meanings of the term 'security' has risen. Traditionally, security has been considered in the e-science arena as being related to user identity and rights assessment which is well covered by this section. But there is another meaning for 'security' that has been typically used in the National Research and Education Network (NREN) arena, which is that dealing with security incident handling. As e-science infrastructures grow and use become more mature, security incident handling coordination becomes critical. It may be worth starting a separate topic on security management in the next e-IRG cycle.

5.3. CONTEXT

Throughout the past months, the growing awareness about the relevance of a global AA infrastructure for e-science has driven the efforts in Europe and worldwide. There are very promising results in the following aspects:

- The consolidation of the regional Policy Management Authorities (PMAs) in the Americas and the Asia-Pacific region, following the model of the EUGridPMA, and the growth and maturity of the EUGridPMA.
- The consolidation of International Grid Trust Federation (IGTF) as the global forum to coordinate identity policies in the global e-science context. This consolidation has translated into proposals for new and more advanced services, including on-line verification services and formal mechanisms to express naming policies.
- The architecture and profiles of the eduGAIN con-federation system within GÉANT2 will provide a distributed and dynamic trust fabric to support the exchange of identity information among different AA infrastructures, possibly based on disparate technologies.
- The advance of the eduroam network access federation worldwide. NRENs of most European countries have already joined the federation, as well as the Australian NREN. There have been reported initial eduroam setups in China, Japan and the US. Some successful experiments using IGTF-recognised identities to access eduroam have been recently carried out.
- The growth of TERENA Academic Certification Auhtority Repository (TACAR) usage and the work towards a new, more global-oriented, version of the repository. Several additional uses have been proposed for TACAR as a neutral point for trust aggregation.
- The emergence of the concept that local identities, either within the context of institutional identity systems or national AA systems, can and should be linked to global identities (such as user certificates in the Grid environment).
  First steps in that regard have been undertaken by investigating short-lived credential services within the EUGRIDPMA / IGTF.

However, there is still work ahead to achieve full AA harmonisation, allow the interoperation of different AA infrastructures and scale them up to reach the whole of the potential user community. Areas that require attention include:

- Dealing with the outreach of the AA infrastructures. This raises questions of scalability, the consensus on a general confederation model, and the generalisation of the use of these infrastructures.
- Dealing with the 'inreach' of the infrastructures, both by incorporating the campus infrastructures in the loop and by educating users in the appropriate management of identity data.
- Dealing with legal issues, mainly related to privacy and liability aspects. More specifically, the 'legal impedance' adaptation between national and regional regulations and practices is a significant risk.

The IGTF envisages the following trends:

- a larger number of identity providers, possibly exposed to the Grid relying parties by proxy authorities;
- A wider variety in the 'quality' of the user-held private data used to protect their identity credentials (i.e. by the use of USB hardware tokens, or integrated smart cards issued by either an institution or by the national or regional government);
- Coupling and clustering of identity providers via bridging technology, in the latter case likely to be supported by a single proxy authority to hide complexity.

5.4. PROPOSED APPROACH

From a purely technological point of view, there are remarkable advances towards enabling a global AA infrastructure, and they seem to have enough momentum to achieve their goals.

But there is a strong need for coordination, both at the European and global levels, to let information on achievements and needs flow in all directions. Initiatives like:

- IGTF itself and the Certificate Authority Operations (CAOPS) Working Group in the Global Grid Forum (GGF);
- the e-concertation workshops and follow-on activities;
- the federation coordination fora around TERENA's Task Forces on European Middleware Coordination and Collaboration (TF-EMC2) and Mobility (TF-Mobility), and Internet2's Middleware Architecture Committee for Education (MACE): Research and Education Federations (REFEDS), Shib-enable, eduroam Global Working Group;
- workshops between NRENs and Grid infrastructure operators (to promote cooperation between these important service providers of the European e-Infrastructure);

play a key role in this coordination process.

A further element which could be tackled in a future edition of the White Paper is the interconnection between AA infrastructure- and Grids-identities.

5.5. RECOMMENDATION

- Trans-disciplinary (Grid projects, NRENs, other user communities) and trans-continental forums that move towards the establishment of a global, seamless AA infrastructure for e-science applications should be encouraged. The e-IRG wishes to acknowledge the efforts made in this direction by the IGTF and the open information exchange point provided by TERENA task forces.

- As stated in the introduction, and as an additional recommendation, the e-IRG would like to introduce a new topic on security incident handling and response coordination in the next e-IRG cycle.

6. RESEARCH NETWORKING POLICY CHALLENGES FOR E-INFRASTRUCTURES

The efficient provisioning of network resources is a key component for ensuring the wide usage and seamless operation of distributed Grid infrastructures. The Grid middleware is evolving so as to integrate network resources and, at the same time, the research networking community is focusing on tools for providing to end-users and applications the ability to dynamically control the allocation of network resources. As technologies and tools for the provisioning of network resources to Grids are evolving, it becomes evident that the limitations are no longer of a technical but rather of an organisational nature. What seems to be missing is a common language between the worlds of networks and Grids and a framework for integrating network and Grid resources on an end-to-end basis, across multiple interconnected domains.

6.1. GOAL

The goal here is to give proper answers to the mostly policy-oriented questions related to the key issues wrt. Research Networks (RN):

- What is the role and the function of RN in e-Infrastructures ?
- How the overall policy wrt. RN should look like?
- How to take into consideration all subjective and objective aspects in the RN policy?
- Why and how demand and supply should and could meet in case of RN?
- What are the conditions of sustainability of RN and the network based services?
- How to finance the RN development and operation?
- When and how to introduce new generations of the RN?
- Why and how to plan for RN development?
- What are the diverse aspects in RN policy and how to take them into consideration?
- Why and how to use dynamically reconfigurable e2e connections?
- How to handle data/knowledge base and d-library, e-repository, e-archive etc. issues?
- What RN policy to apply wrt. education and training in a wider sense?

6.2. CONTEXT

1. Role and function of Research Networking in e-Infrastructures

The frequently asked question here is if the RN is just one component of the resources or should it be considered a crucial element of the e-Infrastructure?

The suggested response is as follows:

Research Networking policy should be based on the postulate that the Network is a dominant major element of the e-Infrastructure. No remote accessibility of distributed resources (computing, storage, data, equipment, and other resources) is possible without an appropriate network infrastructure. Therefore RN policy is a key issue in building the ERA in general, and also in the introduction of demanding e-Infrastructure applications such as complex e-Science activities.

2. RN policy in general

The valid question here is: Why RN policy is needed, why not just use the available data communication means provided by the RN (at present GÉANT-2)?

Suggested response:

The strategy for the development and operation of the RN should be based on principles stemming from an RN policy – that's why a clear RN policy is needed. From the strategy it follows how fulfilling the diverse needs for network services can best be achieved. Obviously, adequate policy should allow perfect service both for less demanding everyday users and most demanding Grid users.

3. Subjective and objective aspects in RN

Coherent and realistic RN policy should assume careful handling of not only objective (technological, organisational, financial, etc) but also subjective (responsibilities, governance, etc.) aspects of RN operation and usage. But how to take into account these complex aspects?

Suggested approach:

RN policy should be based on well defined, mutually accepted principles and practice of how to share responsibilities, duties wrt. governance and control between the RN service providers and the RN users (including the Grid user community).

As far as the details are concerned, it is to be taken into account that (1) governance in RN has been experiencing a proven, stable model (based on the NREN Consortium and on an operational organisation, DANTE) for more than 10 years, and (2) the less demanding everyday users don't need special handling of their service requirements. Therefore the assumption of having a similar governance model on behalf of the demanding users' community (practically the Grid community) leads to an efficient forum, based on a similar governance structure of the parties, for negotiating about the subjective aspects of RN usage. The goal is obviously mutual informing and advising each other, but not intervening into each other's territory – as far as the overall networker-user interaction is concerned.

4. Demand and supply

Agreed RN usage policy isn't possible without well matching service demand and supply scenarios. However, both too much overestimated demands/requests and too much cautious supply/offers may sometimes result in a mismatch between seervice demand and supply. The question is: how demand and supply can best be helped to meet?

Suggested answer:

RN policy should most seriously take into consideration user demands, provided that user requests take into account (1) realistic user needs and (2) realities of both network development and provision of services.

However, there are also some conditions to be met by the Research Network developer/operator. RN policy should be based on (1) investing maximum efforts to fulfil user needs, and (2) interoperability of regional, national, and local segments of the network, and also (3) an integrated, automated, flexible handling of user requests, on the basis of a permanent service provision model.

5. Sustainablilty of services – sustainability of the RN

Permanent service provision assumes sustainable network operation. This means that RN policy should be based on a sustainable network infrastructure. But how long range sustainability of the network can be achieved, especially with respect to the financial conditions of that sustainability? (It is to be taken into consideration that besides national funding difficulties/instabilities also project like funding mechanism for research networking on behalf of the EC may cause uncertainties wrt. sustainability.)

Suggested solution:

An ”old” idea from the early period of planning for the ERA provides a perfect solution for achieving sustainable financing (see http://europa.eu.int/comm/research/era/pdf/networks.pdf) The key is a continuous financial contribution on behalf of the EC, rather than repeatedly running finite lifetime EU FP projects in RN. (This model could later be followed in funding the development and operation of the European ”neutral” Grid – the single, integrated European grid infrastructure involving the single national ”Grids for all”, serving all regions, all academic and research communities, and all disciplines.)

However, also the individual NRENs should do their best for making their networks sustainable (political/legal environment, technical/infrastructural background, stable NREN organisation/operation, technical/administrative expertise, on top of stable funding). Such an overall sustainability, met by the above outlined sustainable EC funding, leads also to financial sustainability not only on the European but also on the nationa level

It is also to be noted that the Pan-European task of achieving sustainability of the RN should be connected with step by step easing the digital divide – making the RN (and later also the Grid) in Europe and in each country the sustainable constituent(s) of the e-Infrastructure.

6. Financing of RN development and operation

Financing the Network (development and operation, including both IP and e2e services) assumes considerable moneys. However, national resources are sometimes, at some NRENs, limited, while annually available moneys from FP7 for RN are, in Q2/2006, still not exactly known. (An important aspect here is that, practically in each European country, the level of national support for RN development and operation does highly depend on the level of the EC funding available for the very NREN.

The suggestion:

Keeping all the above aspects in mind, the suggestion wrt. financing the RN activities is clearly to be based on stressing that the RN policy should definitely aim at fast development based on elevated EC funds.

Multiple reasons for a considerably increased annual EC contribution for RN (especially for the continuation of building, developing, and operating the actual generation of the GÉANT network) have already been tabled not only within the e-IRG but also within TERENA and within the ENPG. The most important such reasons are listed below:

- general political importance of European and global research networking increases
- seamless pan-European end-to-end infrastructure to be developed
- geographical coverage of the pan-European network to be widened
- digital divide to be narrowed
- funding for testing, piloting and introducing new technologies/services needed
- development of access and campus networks to be accelerated
- development of global connectivity to be continued
- dependence of national funding on the level of EC support to be taken into account
- cohesive role of EU funding in the field or RN to be seriously considered.

The level of funding RN in FP6 has been 0.53 % of the entire FP6 budget (93 M€ for the 4 years of FP6). By taking into account the foreseeable 54 B€ budget of FP7, a similar 0.53 % would mean about 280-290 M€ FP7 funding for the development and the operation of the European research network (ie. for the successor of GÉANT-2) – some 70-75 % increase as far as the annual average EU funding is concerned. This is the level the RN community is looking forward from the start of FP7.

However, the NRENs should cover their matching contribution – in average, an equivalent to the EU funding. Here the question arises if it is reasonable, acceptable, and desirable, to change the rule about the 50 % balance between EU funding and national funding? The basic question here is if all the NRENs will be able to cope with the increase of the EU funding if the 50 % rule is kept?

Anyway, considerable flexibility is needed wrt. that 50 % rule since the EU funds definitely should contribute to a more homogeneous European RN, characterised by much less digital divide as in early 2006.

Indeed a wise mix of subsidiarity and solidarity leads to best results.

7. Introducing new generations of the research network

Building a new generation of the RN is time consuming and tedious. It takes at least 2 years from the start (it has taken 2,5 years in case of GÉANT-1 and some 3,5 years in case of GÉANT-2) – independently of whether technology, architecture, topology are fixed or there is a flexibility for adaptive changes.

Therefore the following suggestion is quite obvious:

RN development policy should get independent of the periodicity stemming from the duration of the related EC projects (4-4 years in case of both GÉANT-1 and GÉANT-2). Revolutionary and consolidation phases in RN development/operation should rather be matched to the emergence of promising/proven new solutions – by taking also into consideration that (1) smooth transition is always a must and (2) keeping leading edge position in global scale is a major requirement.

8. Planning for RN development

Planning quantum developments of the European RN should be made in a timely manner, on the basis of the available preparatory documents. SERENATE (a project having run parallel to GN1) published its final report, summarising important suggestions, later than the start of the GN2 project of building GÉANT-2. EARNEST (an activity within GN2) should provide suggestions in time for ”GÉANT-3” (the foreseeable successor of GÉANT-2) - at least as far as the timing of ”GN3” (the foreseeable successor of the GN2 project) is concerned.

The suggestion here is rather simple:

RN policy should carefully take into account the required timeliness of planning RN developments and providing input for preparing plans. Specific activities within GN2 (Joint Research Activities as well as Networking Activities, including EARNEST) do well contribute to such preparations.

Therefore the various user communities of the RN (especially representatives of the demanding users) are urged to participate in EARNEST (and its future successors) aiming at providing input to planning for mid-range and long-range RN development – for the benefit of every users, from the least motivated to the most demanding ones..

9. Diverse aspects of RN

Research Networking faces numerous challenges and the many aspects sometimes contradict each other.

This fact is reflected by the following suggestion:

RN policy should wisely take into consideration a number of different, sometimes contradictory aspects. Although the detailed responses to the diverse challenges in most cases aren't policy level matters, at least the width of the spectrum of the many aspects is to be well taken into account also on the policy level.

Some crucial aspects to be taken into consideration by the policy are listed below:

- Optical networking options (BAU vs. IRU vs. DIY – cable vs. fibre vs. λ)
- Network architecture (traditional hierarchy vs. CBF)
- Multi-domain management/interoperability (standards vs. interfaces)
- AAAI options (NRENs vs. user communities – service vs. middleware)
- Covered user communities (low vs. high end; non-profit vs. industry)
- Areas of activities (networking for R&E vs. R&D on networking)
- Business/economy (returning investments vs. minimal current costs)
- Global connectivity and related business models (cost-sharing vs. accounting)

10. Dynamically (re)configurable e2e connections

Lack of simple, straightforward, direct way of formulating and automated handling of reconfigurable VOPN requests by network users and providers, respectively, is an issue to be duly dealt with in the short range on the RN policy level. But how?

The basic elements of a suggestion with this aspect of making the RN policy are as follows:

Development of policy elements (determining the principles, tools and methodologies) for network resource allocation and VOPN configuration is necessary, as a fundamental part of the RN policy. The involvement of the representatives of the RN community and the major user communities (including the Grid community) is an absolute necessity.

In case of Grid applications, the policy should be based on the need for efficient provisioning of network resources in order to ensure seamless Grid operation. The need for a common language between the Network and Grid communities and for a framework for integrating network and Grid resources on an end-to-end basis, across multiple interconnected domains, is also to be taken into account by the policy. All participating network domains should follow the related policy elements, and also the basic set of principles, together with the suggested methodologies for the establishment of end-to-end connectivity and the allocation of end-to-end capacities.

In accordance with the above requirements, involving the NREN Consortium and a number of user community/consortium representatives in the development process is a must. Paying special attention to the Grid user community, as well as to the crucial QoS, AUP, and accounting (business model) aspects is unavoidable.

11. Handling data/knowledge base and d-library, e-repository, e-archive etc. issues

Issues related to developing and using data/knowledge bases, d-libraries, e-repositories, e-archives are fast emerging and the question does arise on how to handle these issues by the RN policy.

The suggestion here is the following:

The RN policy should carefully be extended to cover the fast emerging issues related to data/knowledge base, d-library, e-repository, e-archive developments and applications, especially to those providing information about the e-Infrastructure resources as well as about the use of the Research Network and also about the user demands.

Here the information for the demanding users about how to request VOPNs, how to use e2e connections, what conditions of the due usage apply, etc. are of key importance. All these background supports should provide necessary knowledge about the network (by an electronically accessible way). The knowledge base should also be appropriate for integration with the knowledge base about Grid parameters and usage information.

The stored, easily accessible integrated knowledge about the Network and the Grid infrastructure should serve as a comprehensive resource of information both for the demanding users and the novices in e-Science applications. Moreover, training and education of the same integrated knowledge about Network and Grid usage should provide an additional opportunity to the actual and potential users for reciving proper, adequate and up to date information on network service models, capabilities, provisioning procedures, costs, etc.

12. RN policy wrt. education and training in the widest sense

The Research Network provides leading edge infrastructure for electronic communication, information access, and remote collaboration. The question here arises: why just the research and higher education communities may enjoy the benefits of the most advanced networking infrastructure?

The suggested response with respect to the possible RN policies is somewhat hesitative here.

Indeed, the RN infrastructure perfectly suits the requirements of seamless, secure, high speed access to the knowledge base of the society (which is also the Lisbon Objective here), not just for research and higher education communities but also, among others, for secondary and primary schools. In fact, accessibility of the RN is extended in a number of countries to serving secondary and primary educational needs, together with similar extensions toward public health organisations, libraries, and other public sector areas of advanced applications.

Involving the objective of following such an extension strategy towards the secondary and primary education levels (ie. following the so called Schoolnet approach) may be an RN policy element on the European level, and certainly may be an element of the RN policy at any NRENs on the national level. However, special care is to be taken to the possibly emerging issues before deeply entering the related development and organisational steps because, although the NRENs are enablers of introducing and extending the Information Society to the young generations, there are considerable differences between the academic users and the teenager users of the network. Therefore, entering the Schoolnet arena by the NRENs needs special attention and special preparations.

6.3. PROPOSED APPROACH

- It is to be kept in mind that Research Networking policy challenges are multifold.
- Careful investigation of the policy challenges is necessary for achieving a good, well applicable RN policy.
- Distinction is to be made wrt. policy coverage. Different user community segments require different handling and somewhat different policy features.
- Above, the general policy aspects have been summarised, with some more emphasis on the demanding (high-end) users. The above aspects are to be taken into consideration in case of determining the general policy, while some further elaboration of the challenges is needed for going towards specific policy versions.
- The findings are equivalently useful for the RN community, for the user communities, for the national funding bodies, and also for the EC. Best results can be foreseen if all these target audiences try to do their best in following the above suggestions.
- Further analysis may help in developing fine tuned RN policies adapted to specific circumstances/conditions.

- The policy should provide a solid basis for straightforward progress toward:
    - well established strategies,
- detailed development plans/steps (in a wide sense), and
- reliable and dependable high quality operation and service provision practices.

6.4. RECOMMENDATIONS

The recommendations are directly stemming from the suggestions appearing in the Context section. Here just a short summary of them is provided.

A. Policy requirements

A1. Research Networking policy should be based on the postulate that the Network is a dominant major element of the e-Infrastructure, supposed to provide perfect service both for less demanding everyday users and most demanding Grid users.

A2. RN policy should be based on well defined, mutually accepted principles and practice of sharing responsibilities and duties wrt. governance and control between the RN service providers and the RN users. The proven RN governance model is available also for RN usage based on realistic user demands and on a permanent service provision model.

B. Sustainability and development issues

B1. Continuous EC financing rather than running finite lifetime EU FP projects in RN would be preferred, together with sustainable operation of NRENs and step by step easing the digital divide. Elevated EC funds (>0,5 % of the entire FP7 budget) in 2007-2013 and changing the 50 % rule are most desirable.

B2. RN development policy should be matched to the emergence of promising/proven new solutions rather than to the periodicity stemming from the duration of the related EC projects. Smooth transitions and keeping leading edge position in global scale are major requirements. Planning for such developments needs also user inputs, eg. through the EARNEST activities.

C. Requirements on network features and on user information

C1. RN policy should wisely take into consideration a number of different aspects: technology options, network architectures, management/interoperability, AAAI options, user communities, areas of activities, business/economy, global connectivity, as well as efficient allocation/integration of e-Infrastructure resources, common language of networkers/users, common practice of network domains, Grid aspects, QoS, AUP, etc.

C2. Informing, training, and educating the users should provide necessary knowledge about the RN and also about the Grid infrastructure. Easily accessible data/knowledge-bases, d-libraries, e-repositories, e-archives should serve as resources of the requested information. Education and training, in general, should carefully consider extending the RN to secondary/ primary schools, by taking into account the differences between academic and teenager users.

7. INCENTIVES FOR RESOURCE PROVISION IN SCIENTIFIC GRIDS

7.1. POLICY AREA

The existence of suitable incentives for resource provision and sharing across resource centres is essential for the efficient operation of scientific Grids. These incentives should operate across multiple administrative domains, and both within and between virtual organisations. Otherwise, in the extreme case, the rational strategy of all peer organizations would be to free-ride on others' efforts - that is, to share no (or limited) resources while consuming as much as possible.

7.2. GOAL

Our goal is to explore the use of realistic incentive mechanisms for resource provision in this context, taking into account issues such as implementation and enforcement, economic efficiency, fairness and also political considerations.

7.3. CONTEXT

The resources exchanged in a Grid system are mainly commodities (e.g. CPU cycles or storage space) and they are rivalrous because when they are used by a user they are not available to others. This fact motivated many researchers to consider the implementation of market mechanisms for determining prices of resources, envisioning a global Grid market either for achieving efficiency and increasing cooperation in resource sharing communities (see HP's implemented market-based resource allocation system Tycoon or the work of Buya et. al) or for financial exploitation (mainly work from Kenyon et al).

However, Grid computing could also, in principle, be operated on a peer-to-peer (P2P) basis with simple contribution and consumption rules (no markets). There are those who envision a global universal Grid (realising Sun's famous statement 'The network is the computer') where all users could share their unused resources to build a single supercomputer for running computational intensive applications. However, such a Grid application is not yet widely deployed on the Internet and some have expressed doubts as to whether this will ever happen. There are, however, more constrained systems under development, limited in the exchange of resources between well-defined and symmetric communities. These include scientific Grids or PlanetLab, in which case more general work on peer-to-peer economics is highly relevant. The most popular incentive mechanisms to enable resource trading between peers (studied mainly in the context of file sharing) are those that are based on tokens equating consumption and contribution of resources or reputation (peers that are positively rated for their reliability, contribution, etc. receive better quality of service).

We explore the issue of resource provision in scientific Grids from a different perspective to that which is usually followed in the literature. More specifically, we focus on the efficient provision of resources (rather than their allocation) and we model resource provision in scientific Grids as a public goods problem. It should be noted that unlike classic public goods, the resources involved in a computational Grid are rivalrous. However, when there are infrequent resource contention overlaps it could be assumed that, over time, Grid resources are non-rivalrous. In this case, the sum of all the resources contributed by peers in a Grid system, which typically consists of the initial investment in infrastructure from each organization, the total 'Grid capacity', exhibits an aspect of a public good. Under a fair scheduling scheme (when resource usage could be scheduled in some non-overlapping way) this public good would then be potentially available to all participants for executing computationally intense tasks. So, an interesting mathematical abstraction of a scientific Grid is a server with capacity C, which is to be provisioned by the clients themselves. The delay experienced by clients certainly depends on the overall rate of requests but the higher the total capacity the less delay they will all experience.

Under this modelling approach there are interesting recent results in the context of P2P file sharing. These suggest that a simple fixed contribution mechanism (in terms of infrastructure invested by participating organizations) can achieve asymptotically the maximum efficiency that could be possibly achieved by any mechanism under the constraints of incentive compatibility and budget-balance.

However, many Grid-specific issues still need to be addressed. More specifically, we should: explore whether the main assumptions made are realistic; devise ways to extend the proposed mechanism when they are not; study many challenging implementation issues related to scheduling, accounting, and enforcement; and examine to what extent these affect the applicability of our theoretical work in the context of file sharing.

Then the central question a system designer should answer is how much each participating organization should contribute towards building a system with an efficient level of capacity C. This is a challenging resource provision problem having many similarities with the classical mechanism design problem in economics known as 'private provision of a public good'. According to the theory, one should take into account the preferences of the different organizations in order to achieve an economically efficient equilibrium. A system designer having complete information about the utilities and costs of the different organizations for Grid capacity could easily compute the provision rule which would maximize the overall efficiency: that for which the total utility for Grid capacity built minus the total cost is maximized. But will organizations have the incentive to reveal truthfully their private information and contribute and consume the efficient amount of resources as dictated by the underlying economic model? The answer is no, since according to the corresponding maximization problem, agents with high utility would need to provide more resources.

This is a standard result of economic theory. A mechanism deciding on a resource provision and allocation rule for a public good based on the declarations of the agents of their private information cannot be at the same time efficient (maximize social welfare), incentive compatible (induce truthful reporting of the agent's private information), and budget balanced (contributions gathered actually cover the provision cost of the efficient amount of resources).

7.4. PROPOSED APPROACH

However, recent asymptotics analysis indicates that a very simple fixed contribution scheme, which requires each participant to contribute a predefined amount of resources, can lead the system to an asymptotically optimal equilibrium in terms of social welfare. This is a general result for the problem of private provision of a public good, which provides theoretical evidence that such a simple and intuitive incentive mechanism is enough to provide the suitable incentives for contribution, at least as far as the public good aspect of the system is concerned. But in order to actually constitute a useful tool for deciding on several important parameters of the mechanism (e.g. the amount of resources required, scheduling policies, etc.) and for eventually designing a candidate incentive mechanism, several Grid-specific aspects should be addressed.

An important characteristic of scientific Grids (and in P2P systems in general) is that they comprise a highly heterogeneous mixture of peers in terms of utilities and capabilities. It is easy to see that the ability to categorize peers into different groups with different fixed contributions and possibly granting differentiated levels of participation (or quality of service) to each group can increase significantly the overall efficiency achieved. Most importantly, we can show that even a very small number of such groups could be enough. Also note that the fact that participation in scientific Grids is not anonymous and there are observable characteristics of the participating organizations (e.g. their size) makes such a categorization more realistic than it is for other p2p systems such as file sharing.

Additionally, such an approach could also eliminate (or constrain) the need for exclusions, which are in many cases undesirable in this context. More specifically, there are situations where exclusions are undesirable for political reasons. For example, certain organizations cannot be excluded from a national Grid even if they have not enough utility or resources to contribute the required minimum contribution. Notably, the administrative costs and the overhead for aggregating very small resource contributions are more costly than useful and it would be actually beneficial if organizations with limited capabilities participated with zero contribution. A suitable mechanism should therefore be built which would ensure that the incentives remain aligned and no organisation can exploit this fact for its own benefit. That is, organizations who will choose to participate as free-riders should receive an appropriately reduced quality of service or have access to a limited percentage of the overall resources provisioned.

Finally, as in the case of file sharing where there is a limited number of peers who offer a significant amount of resources with mainly altruistic incentives, there are expected to be participants in scientific Grids (higher tier resource centres) which will act mainly as resource contributors rather than clients and which will actually build the initial amount of the Grid capacity without any explicit incentive mechanisms in place for doing so. This capacity will then attract a number of user organisations which should ideally contribute to its further increase in order for all to take advantage of the existing externalities and keep increasing the overall utility as the number of participating organisations increases. Thus, even if this initial amount is significant, requiring an additional relatively small contribution by each participant could play a decisive role on the overall efficiency and participation achieved. To this end, the development of the suitable tools for efficiently aggregating all these relatively small resource contributions towards building the total Grid capacity is critical.

The second important aspect of a Grid system is that resources are actually rivalrous: when demand is high, there could be contention for computing power and/or storage. Thus, in this case there will be congestion effects, since the utility of organizations will be affected by the consumption of others (there will be negative externalities). The possible negative effects of contention are magnified by the fact that in reality Grid capacity is not fully available to everybody (as assumed by our abstraction of a Grid system) due to the restrictions that arise by the fragmentation of tasks to different sites. So, scheduling becomes a very important dimension of a candidate economic model, and the hence the corresponding incentive mechanisms proposed should not be considered an independent system functionality. This is, we believe, a very interesting direction for future work in this area.

7.5. WEAKNESSES

Although the public good abstraction of a scientific Grid has some very nice properties and could, ideally, lead to analytical solutions for the efficient amount of resources that should be contributed by participating organisations it is not obvious whether and to what extent the main assumptions made are indeed valid. Towards this end, there was a very interesting discussion at the e-IRG workshop where many Grid-specific issues were raised criticising the applicability of this approach in a realistic scenario.

For example, there are many real situations where the resources of certain organizations may not be dedicated to the Grid and there could be applications that are time-critical. For the first issue we believe that organisations could actually agree to dedicate their resources when the value they receive from the Grid makes it worth doing so; this will be the case when together with all challenging technical issues, suitable incentives are given to all organisations to contribute their resources. When Quality of Service (QoS) issues are important, orthogonal research must be pursued (accounting, Service Level Agreements, pricing, reputation, etc.). A problem with accounting is that reporting sites do not have an incentive to be truthful and contribution may be difficult to measure especially when benchmarks can not be used. Interestingly, in real existing Grids, like EGEE, utilisation is not 100%, and important factors are that not all applications may run on all sites and that individual organisations may choose the sites where their applications will run (according to QoS).

Finally, there are several dimensions of the problem which are difficult to be captured by a simple economic model such as the one proposed. For example, one should be careful to include, correctly weighted, all the attributes of the cost for infrastructure investment such as the corresponding manpower, the operational costs, and the cost of data storage (note that storage does not fit this model because it is consumed and bartering solutions might be more appropriate in this case) and to take into account possible legislative issues. However, in any case, economic modelling fails in general to capture human behaviour. For example, our public good model assumes that there is no correlation in demand between users which is not always the case. Notably, organizations often contribute to Grids for reasons other than simply to get a public good such as reputation, politics, synchronous collaboration, etc..

But despite its weaknesses, economic modelling is often a useful tool for providing insights in order to understand complicated resource provision and allocation problems and design mechanisms which under certain circumstances could lead the system to better, more efficient, equilibria. So, although all the above criticism is valid it would be very helpful to devise ways to extend the proposed mechanism in order to capture important Grid-specific characteristics, at least when the deviation of the underlying assumptions from reality is not such that the public good approach becomes inappropriate. Moreover, one should study the numerous challenging implementation issues related with scheduling, accounting, and enforcement, and examine to what extent these affect the applicability of the corresponding theoretical work. We should stress that we are aware that there are certain types of Grids that do not match this modelling approach. We believe, however, that there is a certain category of scientific Grids for which the public good abstraction is an interesting way to model the resource provision problem and the corresponding theoretical work could lead to practical and efficient incentive mechanisms in this context (which we should note is still a very challenging and open research question).

7.6. RECOMMENDATION

Allowing free access to grid resources does not make sense (at least from an economics point of view), and thus governments and NGIs should consider putting a minimum threshold in resource contribution of a participating oraganization in order to avoid free-riding. This is also seems to be supported by the observation that very small clusters, due to the administrative and technical overhead involved, may be counter-productive.

8. USAGE POLICIES

8.1. POLICY AREA

Usage Policies are essential to regulate the use of the e-infrastructure and to address legal and operational concerns. The aim is to enable interoperability between the infrastructure and other national/international infrastructures, to define the activities of the virtual organizations, to inform users and to protect the operators of services and owners of resources.

8.2. GOAL

This topic presents the current Grid Acceptable Use Policy (AUP) used in EGEE, Open Science Grid (OSG), LCG and several related Grid infrastructures and discusses whether such a simple single policy can meet the needs foreseen for a general e-Infrastructure. The aim is to have one policy for all users and this chapter presents a proposal to achieve this.

8.3. CONTEXT

Production Grid infrastructures run over the pan-European GÉANT backbone and its connected National Research and Education Networks (NRENs). These networks already have well established AUPs as do the majority of the institutes and organizations offering computing resources to the Grid projects and application communities. Use of the Grid infrastructure and resources is therefore already subject to all of these pre-existing AUPs.

Experience to date has shown that a single Grid usage policy is essential to inform the user during his/her registration with a Grid Virtual Organization (VO) as it is impossible for the user to read, understand and accept the large number of individual NREN and site AUPs. The general Grid AUP must be sufficient to satisfy the many NRENs and Grid sites that the users are aware of their responsibilities and that the sites and networks are able to exclude the user from their resource should he/she be found to be using the resource in an unacceptable manner.

The Luxembourg e-IRG White Paper presented a draft of the Grid AUP and the e-IRG recommendation noted the timely operation of the EGEE/LCG/OSG group working on this policy. It also encouraged the group to complete the work on the AUP as soon as possible.

Since the Luxembourg White Paper, work in the EGEE Joint Security Policy Group with Open Science Grid and other Grid projects, has indeed resulted in a single agreed common AUP, which now has to be accepted by all Grid users. This Grid AUP is currently in use in at least EGEE, OSG, SEE-Grid, EELA, EUMedGrid and national Grids such as GridPP.

This fairly abstract and liberal definition of acceptable use (see below) is further qualified in a VO-specific way. Thus, each VO prepares also its own acceptable use policy, which clearly expresses the aims of the VO, together with any necessary details of the applications and/or data involved. Grid infrastructure management and site/resource managers are then able to evaluate the merits of each VO before deciding whether or not to allow the VO to operate on the Grid, safe in the knowledge that all members of the VO will have agreed to the common AUP at the time they registered with the VO.

The Grid AUP was presented and the following questions were addressed during the e-IRG workshop in Linz, Austria in April 2006.

1. Does the Grid AUP approach work for more general e-infrastructures (supercomputing, NRENs, resource centres, etc)? If not, what needs to change?

   The attendees at the workshop agreed that this seemed a good approach and definitely worthy of continued work with the aim of agreeing a common AUP for all uses. It was noted that GÉANT relies on the NREN AUPs (cf. the Geant AUP). Supercomputing Centres each have their own AUP but could perhaps be treated as a VO. The situation with digital libraries and education, rather than research, needs more work and this is outside the scope of the EGEE Joint Security Policy Group.

2. Do users have to be given pointers to all related Usage Policies (NREN, site, etc.)?

   The workshop agreed that for all sites the scale of the task renders it impossible, but if we need to give pointers to other policies it could perhaps be done at the National Grid Initiative level. It was noted that legal advice was required to check whether we have to give pointers to other policies.

3. Are there any legal issues not yet foreseen? (e.g. does the AUP allow operational monitoring, auditing, accounting, etc.?)

   The VOs require access to accounting data and the Grid infrastructures require access to audit logs and monitoring data. There are clearly potential issues here related to personal data privacy. Item 5 of the Grid AUP is there to inform the user that their personal information will be shared for certain purposes. The EGEE Joint Security Policy Group is working on a policy document for accounting to address these issues, but legal advice is required for a more general approach. The Grid infrastructure also needs to look for potential problems with a new VO's AUP before acceptance on the Grid.

4. What AUP should be used for dynamic, short-lived VOs?

   There was no solution found to this problem. It is related to the question of how, and under what policy, does one user grant access to his/her data to another individual? One suggestion is that a general catch-all VO could be considered, although there would be great difficulties here in describing the activities of all such users. It was agreed that there is a need for better understanding of the possible use cases.

5. Are some foreseen uses of the e-infrastructure incompatible with any existing policies?

   It was noted that medical use has special privacy requirements which probably need additional policies. The current NREN AUPs do not allow commercial exploitation of their networks but all agreed that collaboration between the research community and industry is allowed and encouraged.

8.4. PROPOSED APPROACH

The Grid AUP used by EGEE and the other Grid projects is short enough to quote in full:

By registering with the Virtual Organization (the 'VO') as a GRID user you shall be deemed to accept these conditions of use:

1. You shall only use the GRID to perform work, or transmit or store data consistent with the stated goals and policies of the VO of which you are a member and in compliance with these conditions of use.
2. You shall not use the GRID for any unlawful purpose and not (attempt to) breach or circumvent any GRID administrative or security controls. You shall respect copyright and confidentiality agreements and protect your GRID credentials (e.g. private keys, passwords), sensitive data and files.
3. You shall immediately report any known or suspected security breach or misuse of the GRID or GRID credentials to the incident reporting locations specified by the VO and to the relevant credential issuing authorities.
4. Use of the GRID is at your own risk. There is no guarantee that the GRID will be available at any time or that it will suit any purpose.
5. Logged information, including information provided by you for registration purposes, shall be used for administrative, operational, accounting, monitoring and security purposes only. This information may be disclosed to other organizations anywhere in the world for these purposes. Although efforts are made to maintain confidentiality, no guarantees are given.

6. The Resource Providers, the VOs and the GRID operators are entitled to regulate and terminate access for administrative, operational and security purposes and you shall immediately comply with their instructions.
7. You are liable for the consequences of any violation by you of these conditions of use.

Here is an example VO AUP for a fictional VO:

This acceptable Use Policy applies to all members of the Virtual Organization, hereafter referred to as the 'VO', with reference to use of the Grid infrastructure, hereafter referred to as the 'Grid'. The Collaboration Board owns and gives authority to this policy. The goal of the VO is to validate the software they provide to their users (particle physics, astrophysics applications, biomedical communities and others) twice per year within the Grid environment. This procedure needs to cover a wide range of parameters and physical models and therefore demands significant CPU power. At the same time the VO plans to use the Grid resources on a regular basis to make analysis and studies of their toolkit. Members and Managers of the VO agree to be bound by the Grid Acceptable Use Policy and other relevant Grid Policies, and to use the Grid only in the furtherance of the stated aims of the VO.

There are several remaining issues which need to be addressed, some of which are outside of the scope of the EGEE Joint Security Policy Group.

1. Supercomputing centres should be consulted about whether they can use this AUP, with each centre and their registered users being treated as a 'VO'. If not, what needs to be changed to allow this to happen?
2. Digital libraries should be consulted about whether they too can use this approach and if not, what needs to be changed.
3. Legal advice is required on data privacy issues and whether the AUP needs to explicitly refer to other NREN and NGI policies.
4. The AUP for use of e-Infrastructure by education is an open issue. Can it fit into the VO model used?
5. The questions of if and how SMEs/industry can fit into the VO model is also an open issue.
6. There is a need for better understanding of the possible use cases for exploitation of the e-infrastructure by dynamic short-lived VOs before policy issues can be addressed.

8.5. RECOMMENDATION

- The e-IRG encourages wider use of the proposed Grid AUP to promote interoperability and feedback on cases where the proposed AUP does not work saying what needs to change.
- The e-IRG sees a need to investigate how anonymity and/or pseudonymity can be introduced in the Grid environment to better protect the privacy of the Grid user (important in the medical field).
- The six open issues described above all need to be addressed. In particular, the uses of the e-infrastructure which do not fit the Grid VO model are likely to need a different AUP.

9. EDUCATION AND TRAINING

9.1. POLICY AREA

Education and training in the use of e-infrastructure are necessary to realise its potential. They will benefit from shared best practice, coordination and compatible policies, and access to shared education and training material and support systems.

Best practice may include curricula, teaching methods, experience, teaching material and provision of t-infrastructure (i.e. specialised e-infrastructure for education including software, data and computer systems required to deliver courses and to support teachers and learners). Coordination may include consistent policies for student access, presentation of teaching material, and mutual recognition of educational attainment. Shared material may be curricula, recorded lectures, presentations and exercises, virtual experiments and practical experience of collaboration. Shared training support could include course scheduling and advertising as well as the provision of a shared t-infrastructure to support educators and learners.

We seek to establish a consensus and collaboration in the development and delivery of education and training that will improve the provision and use of European e-infrastructure. The goals will include:

- encouraging recognition of the importance of education and training;
- discovering, understanding and recording the current practice;
- developing and presenting agreed best practice;
- establishing mutual recognition of courses and educational or training attainments;
- identifying the opportunities for sharing educational material and training facilities;
- establishing best practice and policies for contributing those shared resources; and
- establishing best practice and policies for providing access to those resources.

9.2. GOAL

The purpose of developing Education and Training policy is to increase significantly the rate, quality, efficiency and extent of education and training. This should accelerate the exploitation of e-infrastructure and increase the benefits for EU citizens. This in turn requires increased investment in education and training. It requires that we identify how best to develop the educators and trainers, the educational material and infrastructure, and hence our capacity to educate and train. It is anticipated that this will be helped by pooling knowledge, sharing in the development of curricula and course materials, and by sharing t-infrastructure. Students and staff should be able to learn and teach across regional and national boundaries. The resulting skills and qualifications should be accepted throughout the ERA.

Investment in education and training is urgent to achieve the benefits from the EU lead in e-infrastructure. We need to disseminate rapidly the knowledge about exploiting e-infrastructure that has already been built to a much wider workforce and potential user community. We need to enable the EU citizens to make good judgements about the value of e-infrastructure for their education, businesses and their personal well being.

If this is left entirely to incremental adoption by universities, schools and training organisations the opportunity will be lost as the natural processes of innovation in education are slow. It is necessary that the e-IRG stimulates coordinated behaviour and a more rapid response.

9.3. CONTEXT

It is essential to make investments in training and education in order that the full potential of e-infrastructure may be realised for the benefit of EU citizens. Awareness raising, management and user introductions, technical training and the development of concepts and culture through education are all necessary. They will equip the citizen to appreciate the potential and to understand and develop new advances in all disciplines. They will also increase the capability to develop and deliver e-infrastructure and its applications.

The chapter on User Support has some overlapping concerns, though it is primarily concerned with supporting the operation and management of grid infrastructures. It recognises the need for education and training as part of developing and supporting users. Consistent policies and collaboration are needed that develop synergy between education and training on the one hand and user support on the other; as the development of a education and training chapter is a new undertaking, this relationship needs to be developed.

9.4. PROPOSED APPROACH

It is necessary to partition the development of Education and Training Policy into subtasks, and then to either address these sequentially or, if there are enough volunteers, to develop several contemporaneously. We plan to augment the available resources by drawing on the work of other groups, including:

- the Training, Outreach and Education Community Group (TOE-CG) being formed at the Global Grid Forum (GGF);
- the Training Concertation group for EU projects led by Dr Rosa Badia; and
- the International Collaboration to Extend and Advance Grid Education (ICEAGE) Forum.

These broader collaborations will have two effects: they will inform our policies with a broader international experience and enable a European influence on their policies.

The steps towards understanding and policy currently under consideration include the following.

Characteristics of Courses. This will develop standard terms for describing courses, e.g. their length, size, pattern of repeat, locations, target audience, form of delivery and target outcomes. For education, outcomes may include an understanding of the ways in which e-infrastructure may be exploited by any specific discipline, e.g. medicine, geosciences, engineering or economics. For training, the outcomes may include competencies in using, developing applications with, deploying or operating various Grid technologies. In computing science courses, the outcomes may include an understanding of the challenges, architectures and systems properties of e-infrastructures for computer scientists.

Collation of Courses. This will develop a snapshot of information about the current courses underway. A course is given to a cohort of students drawn from a particular target population. It will follow a particular curriculum, defining its educational goals and sequence of topics. It will be presented using educational methods and delivered by particular teachers under the auspices of some organisation. It may lead to particular qualifications denoting specific academic attainments. Some courses are recurrent, such as academic courses, e.g. Masters degrees in e-science and courses in specific disciplines that include Grid education and summer schools. It will also include training courses, such as those sponsored by major projects. The collation will capture enough information about the courses to allow the policy group to draw on it for examples, test cases and best practice.

Mechanisms for Agreeing on Curricula and Standards. The e-IRG itself should not attempt to define curricula suitable to meet particular training or educational targets. However, it should establish a framework in which these can be developed, recorded and shared. These would then form the basis for exchanging and recognising e-infrastructure-related qualifications and attainments.

Mechanisms for Developing and Sharing Best Practice. The e-IRG should establish a process for developing and sharing best practice in training and education. This will mine information from the collated course material and from experts in the field. Initial work to prime this process may need to be undertaken within e-IRG, i.e. by those developing this chapter, and by those working in collaborating activities identified above.

Policy for Collecting and Presenting Education and Training Material. The benefits of collating, sharing and cataloguing such material have been well demonstrated. The material can include curricula, lecture-presentation material, lecture notes (for teachers or learners), recorded lectures, in some cases linked with metadata, slide-based illustrations and animations. The material also includes case studies, tutorials and exercises, and the software and data that will support those exercises. There are issues to resolve concerning access, exploitation, feedback, proper recognition and editorial control. The resolutions adopted must motivate contribution and investment in improving and developing the material. How do we maintain this motivation and also permit others to use the material to run commercially viable courses? Frequently the IPR of the contributed material will not belong to the authors but to their employing institution. Can the institutions be persuaded to contribute IPR in exchange for access to other material? Is it still possible, if the material is used commercially, to persuade them to contribute? How can provision and use be balanced fairly? How can we combine the valuable professional developments in eLearning and education with the dynamically developed, open-source academic contributions?

Policy for Supporting Self-paced Learners. Many people wish to teach themselves through direct access to educational and training material, including web-based courses. In some cases, these may have tutorial support, may provide introduction to peer or expert cohorts, and may provide exercises and progress evaluations. How should such facilities be supported? What are the guidelines determining how and when students may gain access to these resources, including educational material, tutorial help and t-infrastructure supporting exercises and demonstrations? Can these provide easy access across national boundaries? How will provision and use be fairly balanced?

Policy for Providing t-Infrastructure. In order to support learners and teachers using eLearning portals or running exercises, it is necessary to provide t-infrastructure. There will be a wide variety of forms of t-Infrastructure, mimicking different forms of e-Infrastructure and supporting a wide variety of learning experiences from e-Infrastructure deployment and management to application development and use. In general, providing t-infrastructure is expensive: software and data must be developed, systems must be set up and managed, and computer resources must be provided. Due to this high investment and the correlated high peaks of demand from educational use, it is sensible to aportion t-Infrastructure costs across many educational users in geographically dispersed locations. What should be the policies for providing and using such t-Infrastructures? How can we ensure that they are fair, so that exploitative behaviour is avoided and providers and users find it worth sustaining and maintaining these systems? A special requirement is t-Infrastructure that enables learners to experience multi-site, multi-disciplinary or multi-national collaboration.

Policy for Supporting Mobility. Both students and teachers will move about the European Research Area (ERA). How can we arrange that their qualifications and their access to training and education systems are supported as they move between Member States?

Many of the above activities depend on each other. The current ordering is intended to take account of these dependencies. There are many cases where policy must balance convenience and encouragement against the risk of exploitative behaviour. It is necessary to invest carefully in establishing the policies with consensus and practical support so that education and training in the use and provision of e-infrastructure thrives throughout Europe. Open and inclusive arrangements will be crucial. They will seek to set up the kind of positive feedback and international collaborative behaviour that is exhibited in the better open source software projects.

9.5. RECOMMENDATIONS

The e-IRG recommends that:

- investments in education and training should be made in order that the full potential of e-infrastructures may be realised for EU citizens;
- significant investment in education should be directed towards improving understanding as to how e-infrastructure may be exploited effectively;

- investment in training and education should be directed towards improving the ERA's capacity to develop, deploy and manage e-infrastructure and the applications that use e-Infrastructure;
- policies and standards for the coordination, collaborative support, interchange and mutual recognition of e-infrastructure should be developed across the ERA.

10. GRID ECONOMY - ALLOCATION AND ACCOUNTING

So far the incentive for connecting resources to the Grid has been primarily research & development and limited tests of Grid production. It is now urgent that e-infrastructures begin to handle the costs of usage of the production level, large scale resources that are currently being connected to the emerging production Grids. There are costs associated with CPU, storage, bandwidth, software usage as well as with support actions. There is thus an urgent need to develop models for a Grid economy that support the allocation and accounting of these resources and that can be deployed on Grids of varying sizes.

10.1. POLICY AREA

Accounting is one of the current key issues when deploying Grids for production. So far, Grids have primarily supported homogeneous user groups with dedicated resources who have joined forces in larger collaborations to exchange and share resources in an open and free manner. Grids connecting general purpose resources and thus requiring accounting policies, procedures and reporting have been much less common. One of the main reasons for this is the lack of general allocation and accounting policies enabling the development of standardized mechanisms for an economy for the Grid.

10.2. GOAL

The goal of this chapter is to provide recommendations for how controlled allocation of Grid resources should be developed and what mechanisms are needed to account for various kinds of resource usage on the Grid for the European Research Area. The intention is also that the recommendations can serve as starting points for the development of accounting systems for commercial Grids.

10.3. CURRENT PRACTICE

There are several studies of various economic models for computing, and Grid computing in particular. However, very few of these models have been deployed in real infrastructures with real users. There is thus work needed to make the connection between the high level modelling and the low level work on accounting that is currently taking place.

10.3.1. Allocation process

Looking first at the allocation of resources there are several different approaches:

- Bartering agreements (zero-sum exchange agreements): An 'all-parties with all-parties' model clearly would not work in a scalable infrastructure and a bank-based model is needed. Each participating party pays into the bank a set of resources/capacity for use by other parties and in return receives the right to use other parties' resources. Users implicitly make withdrawals when jobs are run on banked resources other than the resources with which the user is implicitly affiliated.
- Allocations committee: Access to centrally owned resources is granted by an allocations committee handing out credits based on assessment of the quality of proposals submitted by groups of users to funding agencies. The users would then themselves decide on where and when to use the credits. A challenge in this model lies in making sure that credits are internationally negotiable.
- Pooling of real funds: This becomes similar to a 'Service Level Agreement with a Corporation' model in which user groups would contract with the corporation for defined services.

At the Linz workshop it was concluded that there will most likely not be one unified allocation mechanism and that the Grid Economy must support a diverse set of allocation mechanisms on local, national and international levels. It is also emphasized that the ownership of the resources will dictate the allocation mechanism.

10.3.2. Accounting

Standard scheduling systems as well as operating systems commonly have built-in accounting systems to track resource usage. However, they often assume a homogeneous run-time environment, and they lack standard and uniform ways to obtain and represent information from several heterogeneous resources. Thus there has been a strong need for Grid accounting systems that integrate local accounting solutions similarly to the way Grid meta-schedulers and co-allocation managers coordinate, and administer job submissions across several schedulers.

Most Grids do already at this point have their own systems for accounting, which is disappointing and indicates the need for a coordinated effort with a foundation in a high level policy. Examples of Grid accounting projects are Distributed Grid Accounting System (DGAS, Italy, EGEE), (SGAS, SweGrid), APEL (EGEE) and work carried out by TeraGrid and Open Science Grid. In these projects, progress on metering has recently been made and there is also scope for exchange of records and aggregation of information from multiple Grids at a higher level, e.g. (APEL/ DGAS), and (OSG/TeraGrid) (SGAS/APEL).

When it comes to accounting, there is a need to develop an e-IRG consensus on:

- what information about users and usage should be collected and how to assure compliance with privacy laws, etc.;
- formats for representation of this information so a non-invasive Grid accounting system can supplement local accounting systems based on open standards (web services);
- access rights to different data collected for different constituents (public, users, agencies, sysadmins, resource owners/providers);
- reporting/dissemination forms for the different constituents;
- normalization/equivalencing of various resources;
- quota enforcement policies.

It can be concluded that accounting/metering is now becoming available on Grids. There is, however, no interface with allocation mechanisms.

10.4. RECOMMENDATION

- The e-IRG considers that scalable resource allocation mechanism(s) must be included in the design phase of any Sustainable European Grid Infrastructure;
- The e-IRG supports the development of a versatile Grid economy model supporting those allocation mechanism(s) for a large number of small VOs, which should influence the design and implementation of any scalable and Sustainable European Grid Infrastructure;
- An accounting policy and standard that enables interoperability must be developed immediately, but be based on the vision of a full Grid Economy.

This work should be based on an information exchange and collaboration between the e-IRG at the policy level, the Global Grid Forum (GGF) for standardisation efforts, and accounting system developers to make sure that implementation considerations are taken into account.

11. MIDDLEWARE INTEROPERABILITY AND REPOSITORIES

11.1. POLICY AREA

Globally there are many e-infrastructure deployments taking place - EGEE, UK National Grid Service, TeraGrid, NAREGI, etc. While many of these deployments have, or are moving towards, middleware that is based around web services, there are still many issues relating to interoperation and interoperability between these different deployments that need to be addressed. In addition the experience from these e-infrastructure deployments is that there is not a 'single' distribution of software that will meet all of their needs. The ability to compose software solutions from multiple providers is vital to meet the broad needs of the e-research community that will build on top of the deployed e-infrastructure. In such an environment being able to assess the quality of software from multiple sources, and to be aware of the interoperability issues between different services, is key.

Recent community activity within the Global Grid Forum is looking to tackle some issues of interoperation through the Grid Interoperation Now (GIN) initiative. This is a pragmatic activity that is looking to see what can be done now. Longer term effort is needed to build interoperability between different middleware implementations through the adoption of open standards, and the identification of middleware that complies with these standards.

11.2. GOALS

There are several technical issues that need to be reflected at a policy level - both to achieve interoperability and to promote improved software quality:

- Service Oriented Architecture: Service Oriented Architectures based around web services are becoming the norm for building distributed collaborative infrastructures. Interoperability between different providers of web services can only be achieved through the adoption of standards - both to define the infrastructure needed to deliver services - and the services themselves.
- Software Repositories: It is currently very hard to find the software outputs from current and past EU projects. Projects should be encouraged to contribute their software outputs into repositories (along the lines of SourceForge[?]) where they can be maintained for other projects to build upon and use. This need not just apply to open source projects but these should certainly be encouraged.
- Standards Compliance and Quality Assurance: Software repositories (such as SourceForge[?]) may contain software of varying quality. Software in the repository needs to be assesed for its standards compliance and its quality. Clear grading of the software will promote its adoption and reuse within other projects allowing global communities to build around the work.
- Distributions: To promote the assembly of the software services and tools contained within the repository into distributions suitable for particular communities. The effectiveness of these distributions, their portability across different platforms and the interoperability of services within the distribution, also need to be verified and graded.

A software repository containing outputs from EU projects, that have been clearly graded for quality and standards compliance, would promote the uptake of this software worldwide.

11.3. CURRENT PRACTICE

Several middleware projects have been already been supporting the use of repositories to curate software. Interoperability is currently being pursued through *interoperability tests* and the need for more formal compliance tests is beginning to emerge.

11.3.1. OMII-UK

The Open Middleware Infrastructure Institute UK (OMII-UK) which was established in January 2001 hosts a repository of middleware packages contributed from its managed programme and contributions from the UK e-science community. Further details can be found at http://www.omii.ac.uk/.

11.3.2. Globus

The Globus Alliance has recently launched an open source development portal that provides a repository of some of the individual software activities taking place within the Globus community. Projects within the portal are supported by a managed development, build and testing infrastructure, and conform to a common licence. This infrastructure is used to manage the production of the Globus distribution, itself a project within the portal. Further details can be found at http://dev.globus.org/wiki/Welcome.

11.3.3. ETICS

The ETICS (eInfrastructure for Testing, Integration and Configuration of Software) project currently provides a mechanism to securely access a build and test infrastructure that will lead to a repository of components that build on a variety of platforms, with well-known and controlled external dependencies. In the longer term, the goal of ETICS is to propose a certification process for Grid and distributed software that provides a quantitative measure of the quality level of the software component. For further details see http://etics.web.cern.ch/etics/.

11.3.4. EGEE

The EGEE (Enable Grids for E-SciencE) project maintains and deploys a software infrastructure for use by scientists and engineers in 9 disciplines including high energy physics, life sciences, earth sciences and computational chemistry. This global e-infrastructure is run on a production basis with 24 hours a day, 7 days a week availability. The gLite software distribution used within the project is assembled from both internal development teams and external providers before being tested and verified. Once its quality has been established on seperate resources it is deployed to the production resources. For further details see http://www.eu-egee.org/.

11.4. PROPOSED APPROACH

The software deployed to instantiate Grids will be drawn from many sources. It is important to understand the behaviour of the software, its quality, and to demonstrate this across different platforms and environments. This is the goal of a major middleware new EU funded activity - OMII-Europe - that intends to address interoperability concerns between different middlewaresoftware distributions by examining their compliance to agreed standards. Its activity in this area will consist of:

- providing a repository of middleware software components (i.e., services and tools) that are demonstrated to build in different environments;
- providing measures of the software quality;
- recording standards compliance and interoperability of the software components.

This work project does not include any significant new software development, but instead it will build upon existing activities taking place globally in global grid middleware development -and integration - from experiences within EGEE, OMII-UK, ETICS, Globus, and the NSF Middleware Initiative Build and Test framework - and will be driven by the software engineering and deployment needs emerging from these and other projects. As such it will look to provide a repository of software, and to report on its quality and interoperability, to aid the community in selecting software for their own deployments.

11.5. RECOMMENDATION

- The e-IRG strongly supports the continued harmonisation of software services and tools contained in Grid middleware distributions through their engagement in standardization activities driven by practical implementation experience. Self-certification of standards compliance is not deemed sufficient, and the e-IRG considers the provision of an independent documented process for standards compliance and quality assurance to be vital.
- The e-IRG considers that the provision of a repository to include software components (with documented quality assessments) from other EU projects is an important area which should be discussed at future meetings.. It would provide a showcase for the EU community, allowing the the software outputs from various projects to be made available from one location.

The currently funded work within OMII-Europe and ETICS is considered to be crucial in this respect.

12. USER SUPPORT

12.1. POLICY AREA

User support, covering a wide range of users and disciplines, is essential to ensure that Grid technology is used effectively. It must support the user through all stages of development, embracing the initial Gridification of applications, the continuing enhancement of the application, and the day-to-day technical support for running on the Grid infrastructure. It is important that we define policies which will be sustainable in the different national environments.

12.2. GOAL AND CONTEXT

To make production Grids a reality for a broad range of disciplines we must make Grids easy to use. This requires high quality professional support for a large range of users who themselves have a wide range of expertise. This support will have to cover the provision of training and high quality documentation, the Gridification of applications and their day-to-day running. All of these areas must cooperate in the provision of a coherent set of user services. In the two previous White Papers, recommendations have been made regarding models of user support. We update these based on the experience of the past year. Note that opinions are largely based on experience gained by the EGEE projects, which is using a federated model for User Support.

12.3. REVIEW OF PROGRESS AND ISSUES ARISING SINCE THE LUXEMBOURG E-IRG WHITE PAPER

Previous White Papers have drawn on the experience of the EGEE project in the design and implementation of a federated model for user support. We continue to draw on this experience here. The EGEE Global Grid User Support (GGUS) system has now been commissioned and supports users from a wide range of disciplines. A brief summary of the current system structure is given in Technical Annexe ‘GGUS’. It provides 24*7 coverage by 3 teams in different time zones. The running so far has validated the federated approach with 9 Regional Operation Centres (ROCs). The user can via his local help-desk interface generate a ticket which is routed by the GGUS software to the appropriate Ticket Processing Manager (TPM), handling tickets for the suspected problem area whether it be middleware, networking or the application. There is limited experience in handling a large number of tickets and currently GGUS has successfully handled 200 tickets per day.

The experience gained in the past year has shown that it is essential for community Virtual Organisations(VOs) to really buy-in to the system by providing support for the VO TPM which has to deal with application specific problems. This has proved to be a limiting factor in dealing with problems, and the VOs must find resources for this.

The next year will see GGUS coming under heavy load, since the number of VOs using the EGEE infrastructure is growing rapidly. Also as the use of Grid technology broadens, the technical background of users will decrease, making it even more essential that the VO communities have their own expertise providing support.

In parallel with the GGUS, day-to-day support experience has been gained with providing support to new areas for the interfacing of their applications to the Grid. An overview of a model for such work is given in Technical Annexe 'GILDA' (Grid INFN Laboratory for Dissemination Activities) which describes practical work established by a dedicated team and a facility for getting applications started.

The GILDA infrastructure, involving an experienced team and 20 sites over 3 continents, has worked directly with more than 10 VOs. There has been direct contact with users via workshops and practical tutorials. The work has been very hands-on, showing the functionalities and limitations of the Grid, and beginning with practical, simple use cases. This is then followed by practical consulting help in porting applications. This is a very attractive and effective way of getting new applications onto the

Grid. It would seem an attractive model for implementation on the national and regional levels with suitable sharing of valuable resources.

The GILDA activities are very closely related to training (a topic dealt with elsewhere in this White Paper), and indeed within EGEE, the GILDA team worked effectively with the team providing training infrastructures.

An interesting question is the evaluation of user satisfaction with Grid technology. Given current experience with questionnaires it seems that one needs to use the VO managers as filters for acquiring information from their own VOs, since each VO has its own profile in expertise, organization and Grid computing problems. It is very important that the community has well defined mechanisms for evaluating user satisfaction.

12.4. FOLLOW UP ON RECOMMENDATIONS FROM THE LUXEMBOURG WHITE PAPER

We now review and comment on the relevant recommendations from the Luxembourg White Paper (which are in italics)

Infrastructure planning should include the provision of resources for the co-ordination and delivery of an educational programme covering all aspects of Grid usage.

- A strong connection between education, training and user support is needed!
- New users need user support after training.
- The coordination of information for Grid users is a challenge to be met.
- We must establish generic connections to middleware providers in order to interact with the suppliers of components.

We recommend that the infrastructure planning includes the provision of resources for the support of application areas, both in the form of support for the first application migration, and with dedicated ongoing support for application areas according to their scope and complexity.

- The resources and input you can expect from a VO in this area depends on the size of the VO.

We recommend that the developments of day-to-day support systems for the Grid be fully supported.

- The federated model of EGEE should be pursued.
- The National Grid Initiatives (NGIs) should contribute to a European Grid User support system.

It is important that user support system developments are fully documented and discussed in the Global Grid Forum (GGF) and with other major Grid projects. This applies also to the development of information systems for Grid users.

- We must pursue this area in GGF to ensure sharing of experience and technology.

A policy for networking support for Grid infrastructures should be agreed and made mandatory for all entities contributing to the operational environment. The implementation of the networking support should be interfaced to the overall user support infrastructure.

- Connections are already established, but more discussion is needed .
- Is there a centralized user support in networking?
- Connection to supercomputing centres/projects?

The e-IRG has reviewed the recommendations for Grid user support in the key areas of user education, the provision of easily accessible user information, support for applications and the day-to-day running of the Grid and networking infrastructure. The e-IRG would like to stress that the *planning for such support of the infrastructure should allow for the continuity of support in national structures beyond the end of major Grid projects.*

- The NGIs should address user support in order to ensure sustainability of e-infrastructures.
- NGIs should be coordinated on European level.

We recommend that the developments of day-to-day support systems for the Grid be fully supported.

- National Research and Education Networks (NRENs) should become more involved in the process to provide day-to-day Grid user support. This is in progess: GGUS and the NRENs are working on a collaboration between network and Grid support infrastructure through an interface between GGUS and the ENOC (EGEE Network Operations Centre).

12.5. *TECHNICAL ANNEX*

12.5.1. GGUS

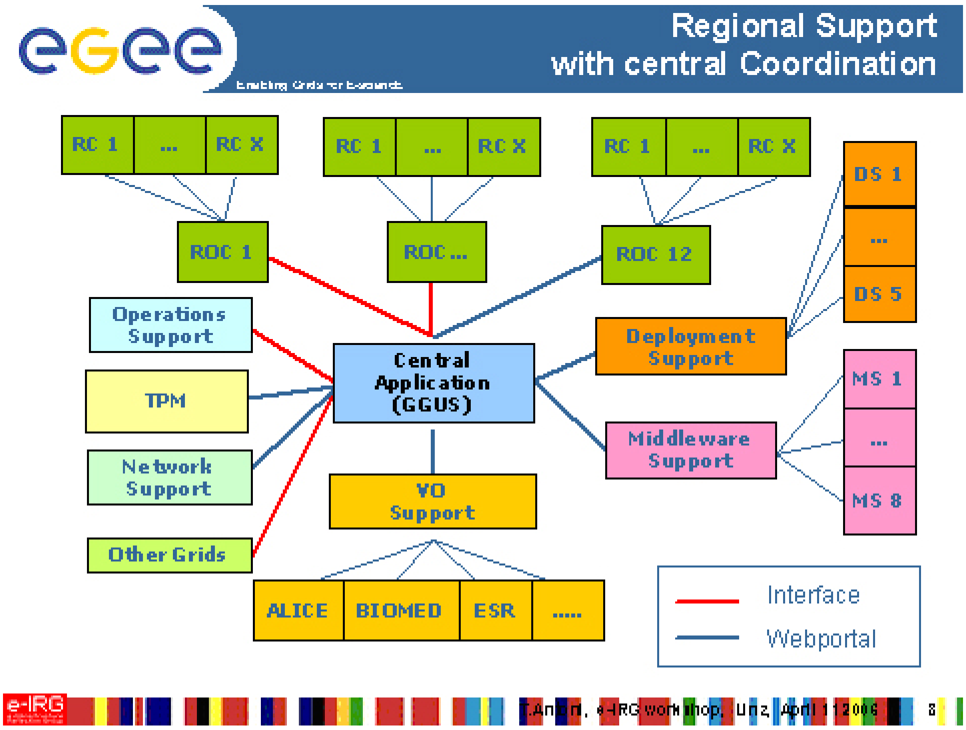

The picture above shows the model of the EGEE user support central application 'GGUS' interfaced to support components for operations, ticket processing, networking, middleware and VO support.

The key points of GGUS are:

- It provides a single entry point for reporting problems and dealing with the Grid.
- It offers a portal where users can find up-to-date documentation, and powerful search engines to find answers to resolved problems and examples.
- Common solutions are stored in the GGUS knowledge database and Wiki pages are compiled for frequent or undocumented problems/features.
- It offers hotlines for users and supporters and a virtual rooms videoconferencing system chat room to make the entire support infrastructure available on-line to users.
- It is interfaced with other Grids' support infrastructures such as Open Science Grid and NorduGrid.

- It is used for daily operations to monitor the Grid and keep it healthy. Specific user problems can be directly communicated to the Grid Operation Centres and broadcasted to the entire Grid community.

12.5.2. GILDA

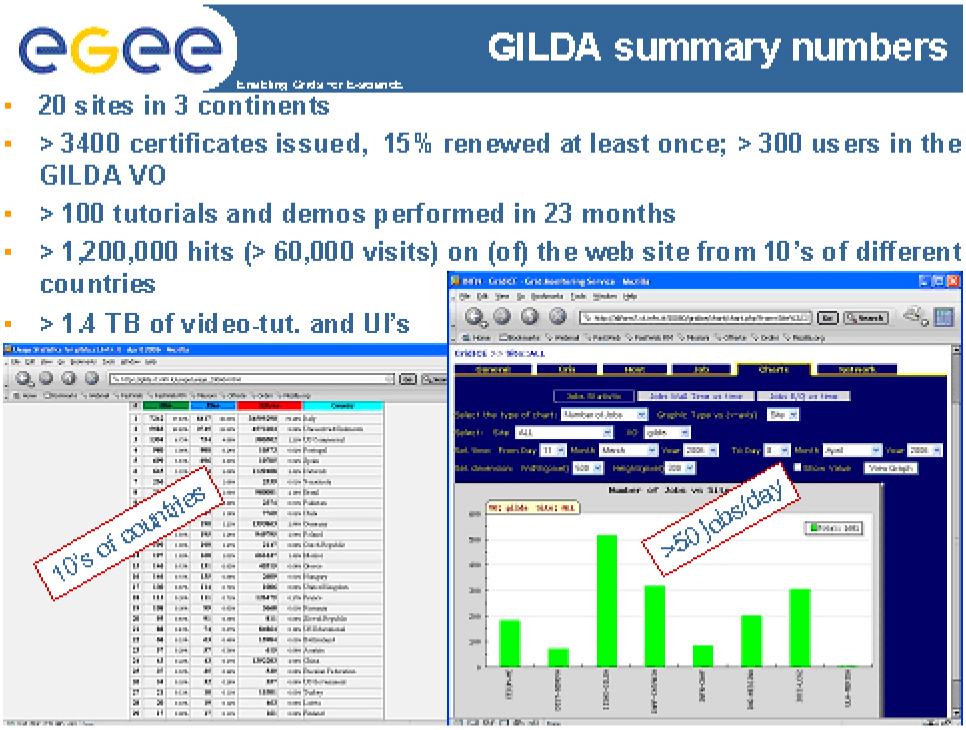

Above we see a summary of the utilisation of the GILDA test bed by many VOs in the past 23 months. The population of the GILDA VO is dynamic by definition, depending on the patterns of courses and workshops. Once VOs are initiated they operate within their own VOs in Grid production infrastructures.

Video tutorials have proved an interesting development in the past year.

13. SUPERCOMPUTING AND PETA SCALE SUPERCOMPUTING HARMONISATION POLICIES

13.1. INTRODUCTION

The present organization of High Performance Computing (HPC) in Europe is based on the integration of leading national supercomputing services. Typically, national supercomputing infrastructures currently operate leading systems that are in the tens of teraflops range, and the natural evolution of peak performance at constant price will lead to a few systems in the hundreds of teraflops range emerging in the next few years. This context does not look sufficient to meet grand computational challenges in research and technology in a way that could match the strong drive of the USA and Japan, where petascale computational systems are planned for as early as 2008.

Europe requires therefore a special action to close the gap and to provide a limited and restricted number of petascale systems, which will complete the present fabric of entry point, midrange and terascale systems available for research and development in the member states. This is a necessity because petascale systems are required by many disciplines in Europe, and there exist many scientific scenarios which provide a strong case for them. It is clear that scientists will adapt their research strategies to the resources available to solve their scientific problems; thus the complete absence of new generation petascale systems may severely handicap European research and development in the next few years, because many new and innovative research subjects and computational applications will simply not be established.

13.2. A TWO-TIER HPC GLOBAL INFRASTRUCTURE

We know today that Grid supercomputing infrastructures like DEISA can add substantial value to existing national supercomputing infrastructures by integrating services and enabling a cooperative operation of a global supercomputing environment. However, Grid infrastructures alone cannot meet all the exceptional requirements of leading, grand challenge HPC applications. Deploying a huge, tightly coupled parallel application on two remote systems is tempting, but if communications are important this strategy (the good old ‘metacomputing’ of the late 1990s) will in many cases be totally inefficient because of the huge latencies (‘delays’ for network practitioners) induced by the finite signal speed propagation. Indeed, MPI latencies are boosted in a continental wide area network from a few microseconds to tens of milliseconds, and this will in general conspire against high overall computational performance of the distributed application. Metacomputing is a good strategy only when wide area network communications are rare.

Many leading grand challenge applications need to be executed in a single system to meet their extreme performance requirements. This is the basic argument in favour of the deployment and operation of a limited number of new, leading petascale systems in Europe. Supercomputing Grid infrastructures like DEISA are not in conflict with this view. On the contrary, they have been designed with future shared petascale systems in mind, because many of the basic Grid supercomputing services being deployed – like the global data management infrastructure – aim at facilitating the access and operation and enhancing the outreach of such future shared petascale systems.

There is therefore emerging consensus in Europe on the necessity of completing the existing Grid supercomputing infrastructures – based on the strong integration of national services – with a restricted number of leading petascale systems. This leads to a ‘two tier’ architecture for the global European HPC services: a Tier One (T1) of a few petascale systems strongly integrated with the Tier Two (T2) of leading national terascale systems in a unique supercomputing environment, with a unique and precise global operational model.

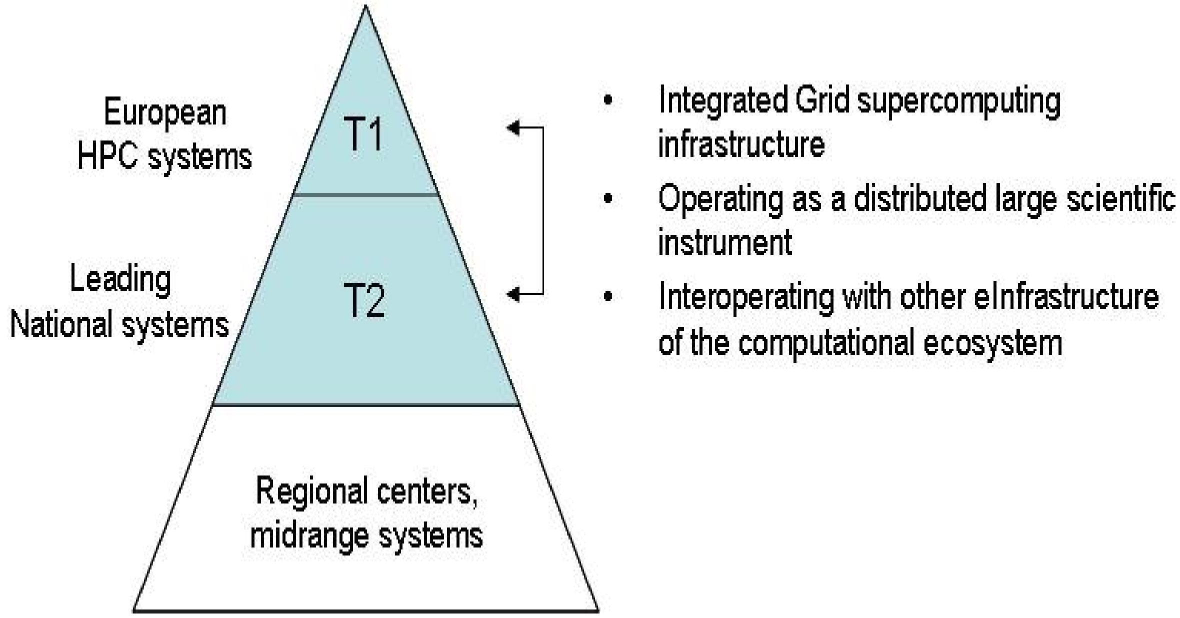

The strong integration of T1 and T2 is required because leading end users will continue to rely mainly on T2 services for most of their work, with access to T1 for the most demanding computations required for their simulations. End users will also continue to access their traditional data repositories, and T1 and T2 systems must be able to cooperate in a global Grid environment. Simulation chains based on workflow architectures must be able to integrate both levels of performance in a unique simulation flow.

The strong integration of T1 and T2 is natural because both tiers are accessed and operated as national scientific instruments like telescopes or particle accelerators, with access based on peer review of scientific projects. This is the way T2 systems are accessed and operated today, and there is little doubt that T1 systems, if and when they exist, will be accessed and operated in the same way. This global infrastructure with very natural boundaries must in turn interoperate with the extant infrastructures of the world wide computational ecosystem.

Another naturally emerging consensus is the priority for deployment and operation of large scale multidisciplinary infrastructures. At this scale, with major national investments required, discipline-oriented computational infrastructures do not seem a necessity. The T1 layer should therefore be multidisciplinary. A few complementary petascale supercomputing architectures available at any one time in Europe would be sufficient to serve the requirements of practically all the scientific disciplines with grand challenge requirements. Discipline-specific services could in some cases be deployed in the T2 or in more peripheral layers.

A balanced relationship between T1 and T2 systems and services is required. A substantial number of member states will maintain a sustained activity at the T2 level, with an occasional access to T1. The policies and funding models to be used for the deployment and operation of T1 must avoid weakening the financial support for T2 on the one hand, and the excessive concentration of HPC leadership and know-how on T1 sites on the other hand. Ideally, T1 sites should of course have extensive computational resources, with services, leadership and know-how remaining distributed over the whole global infrastructure. This is one clear political and strategic benefit of the deployment of a global distributed infrastructure 'à la DEISA' strongly integrating all T1 and T2 systems.

13.3. NEW PROGRAMMING MODELS, NEW ALGORITHMS, NEW METHODS

The e-IRG white paper - Luxembourg 2005 version – states that 'Scientific software is an essential and integral part of the European e-Infrastructure.' Many computer architectures available today are well adapted to specific applications profiles (vector supercomputers are a good old example, but they are not the only ones). As systems grow in complexity (number of processors, number of cores, memory, interconnection network,) special efforts are needed to optimise and tune the applications for efficient operation, with the understanding that not all applications are adapted to run on one specific platform. It is not always possible or reasonable to port existing software to new platforms. In many cases, requirements for specific and special development may become very important. Development of new models and computational methods should go together with the deployment of new, leading petascale systems.

Reinforcement in research of new programming models, new algorithms and new methods is necessary for the better utilization and performance improvement of future petascale supercomputers. Present experience confirms that:

- in some research areas, it is necessary to rewrite important components of the software in order to adapt them to large number of processors;
- in particular, I/O strategies are of critical importance for large applications, and they often require a complete revision of existing codes;
- current algorithms are not able to exploit all available computational resources;
- existing analysis tools needs to evolve in order to be able to help developers in petascale systems.

All of these considerations point to the necessity of a set of services that will complement those of the infrastructure and that will help the scientific community to improve the performance of their applications.

13.4. HIGH PERFORMANCE ACCESS TO REMOTE DATA

In a distributed supercomputing environment, high performance access to remote data is a key ingredient to comply with the HPC requirements. The amount of data a grand challenge application can generate in a petascale system requires very large storage capacity, and, as explained above, large data sets will need to be accessed from remote computing platforms. High performance remote I/O and file transfers become critical services in the presence of shared petascale systems. There are two important aspects in the access to remote data:

- full integration and cooperative operation of the T1 and T2 system;
- full integration of the data repositories with the T1 and T2 systems.

13.5. EUROPEAN HPC POLICIES

The critical issue to be debated is the funding and operational model of the T1 layer. One obvious option is joint capital investments from member countries, with a contribution from the EU. This issue is delicate because of the different levels of involvement in T1 from member countries: not all European nations currently require urgent, sustained and persistent access to T1. Setting a viable funding and operational model is a prerequisite to establishing the scientific policies related to the operation of the T1 layer, and to discussing issues such as:

- who can apply for the usage?
- who decides on the usage?
- what is the cost for the user?

13.6. RECOMMENDATIONS

- The e-IRG recommends that the existing Grid supercomputing infrastructure – be completed based on the strong integration of national services – with a restricted number of leading petascale systems, oriented to a ‘two tier’ architecture for the global European HPC services

- The e-IRG encourages the reinforcement of research into new programming models, new algorithms and new methods that will enable the above-mentioned set of services to be established.

14. RESPONSIVE GRIDS - SHORT JOBS AND POLICY ISSUES

14.1. POLICY AREA

Responsive Grids are defined as Grid infrastructures that provide extensive support for on-demand computing and enable interactive usage. These two features are critical for ensuring that the Grid component of the e-infrastructure will adequately respond not only to the needs of most traditional high-performance computing, but also to the explicit requirements of the full set of e-infrastructure applications. These range from e-science applications, including numerically-oriented ones, to e-commerce. Supporting responsiveness will have a particularly strong impact on the role of the Grid in e-health and e-culture.

14.2. GOAL

In the 1970s, the transition from batch systems to interactive computing was the instigator for a major technological shift. Computers ceased to be specialised tools located in computing centres (be they in businesses or scientific institutes), and became commodity resources. The Grid is facing the same challenge. Large-scale user involvement requires seamless integration of Grid computing power and storage capacity into everyday Information Technology (IT) tools.

Experience shows that an incremental approach in the framework of existing Grid middleware is fully viable. Existing Grids can provide pilot services implementing new policies to enable responsive behaviour: the maturity of this area (including end-user experience) is now sufficient for promoting responsiveness as a concrete first-order requirement for Grid development.

14.3. CONTEXT

14.3.1. Scenarios

Most potential or actual e-applications imply responsiveness: a few examples are data analysis (high-energy physics, medical, etc.), emergency management, financial analysis, collaborative work and dynamic data-driven applications. The common factor for responsiveness is the presence of a human in the computation/interpretation loop. The examples below give three concrete use cases.

- e-health

In a clinical context medical image analysis (segmentation, registration) and exploitation (augmented reality for planning interventions or intra-operative support) requires full interaction, because computer programs cannot yet competitive with the human visual system when it comes to mining these structured and noisy data. Combining the medical user's expertise with the resources of the Grid in compute- and data-intensive tasks is a promising way to transfer experimental research first to clinical practice, and then to routine clinical practice. Analysing large images at a sufficient speed to support smooth visualization requires not only substantial computing power, which can be provided by the Grid, but also unplanned access and sophisticated interaction protocols.

- e-culture

Digital libraries are a key component of e-culture. Most of the resource consumption in digital libraries management is related to en masse off-line tasks such as indexing. When humans query this massive amount of data, various actions are triggered such as feature extraction in a query-by-example scheme, which must take place before the actual search can be carried out, or protection (e.g. watermarking). User satisfaction requires nearly instantaneous response.

- scientific data analysis

In the analysis environment, for example in high-energy physics, multiple users compete for resources and the latency between issuing a task and the availability of the first results becomes important and determines the number of iterations a user can achieve in a working day.

In the first example, a close interaction takes place, the so-called computational steering. The second one exemplifies the case of a visualization/decision loop. This is a very common scheme, where users require only on-line progress monitoring of their results to decide about further actions. A similar scheme is especially frequent in hard e-science (numerical computing, high-energy physics analysis, etc.).

Users of departmental clusters are reluctant to switch to Grid computing (which would provide them with much needed resources), because they need a preliminary phase of interactive exploration of the parameters space. The possibility to easily switch between local resources (for development, debugging and some optimisation) to the Grid (full-statistics runs) is an additional component required for the success of Grid computing. Here responsive services will encourage users to move a larger fraction of their activity to the Grid environment and hence to share their local facilities via Grid mechanisms.

In the larger perspective of ubiquitous computing and ambient intelligence, multimodal interfaces that are capable of natural and seamless interaction with and between individual human users are mandatory. Responsiveness is a key component for Grid-enabling the methods and technologies that form the back-end of these interfaces, such as pattern analysis, statistical modelling and computational learning.

14.3.2. Past experience

The traditional high-end resource providers (parallel computing) have never really implemented time-sharing, despite a significant effort in the 1980s based on virtualisation. Thus, the difficulty of the goal should not be underestimated.

There has been a significant amount of EU-funded research and development targeting responsive Grids, especially the past EU CrossGrid project and the current projects int.eu.grid and GridCC. CrossGrid has developed middleware components and expertise in the fields of scheduling, application monitoring, and Grid access, including support for truly parallel (MPI) execution. This project has demonstrated large-scale and comprehensive applications in the area of flood prediction, high energy physics, air pollution modelling, and medical imaging. GRIDCC targets the most demanding case for responsive grids, which is the remote access to and control of distributed instrumentation.

The EGEE (Enabling Grids for e-Science) project has endorsed the goal of supporting responsive Grids through the creation of the Short Deadline Jobs (SDJ) working group. gLite 3.2 will implement the new services proposed by this group. EGEE has also provided support for Grid-enabling interactive applications in the biomedical area. These activities, originally proposed by the biomedical user communities, have garnered a lot of interest across the other user communities.

The Interactive European Grid (int.eu.grid) will guarantee compatibility and interoperability, but additionally represent a functional extension of EGEE. As well as supporting the execution of EGEE-like batch jobs, int.eu.grid will allow distributed parallel computing with MPI and improved means for collaboration. This will be achieved by exploiting the expertise generated by the EU CrossGrid project to provide researchers with interactive and simultaneous access to large distributed facilities through a friendly interface with powerful visualization. For this reason, int.eu.grid can be seen as a spearhead of EGEE, working in complementary application domains that cannot be fulfilled with the present EGEE middleware. The core participants of the int.eu.grid project also participated in the European pilot project, CrossGrid. Furthermore, several partners of int.eu.grid are participants and stakeholders in EGEE initiative.

On the other hand, many applications faced with the immediate issue of fast turnaround time, and the need for interposition agents (which is related to responsiveness), have chosen to develop autonomous software, which diverges from the main middleware development path of EGEE. Typically, these applications use the basic EGEE infrastructure, but do not share the more advanced submission and scheduling facilities. If this situation were to persist, it would possibly have negative effects on defining and implementing consistent policies for Grid usage, as well as for authentication, authorization and accounting. Furthermore, the unnecessary variety of middleware would not facilitate induction or training.

14.3.3. Technology trends

In industry, a strong trend is emerging toward virtualisation. Virtual machines provide a powerful layer of abstraction in distributed computing environments, providing freedom of scheduling and even migrating an entire Operating System and associated executions. While the relationship of these technologies and the Grid goes far beyond the case of responsive Grids, virtualisation could be a core enabling concept for responsiveness. However, the history of computing suggests the need for some prudence before advancing down this path. Virtualisation has emerged three times: in the 1970s, in the 1980s and now. In the 1970s, its explicit goal was to enable time-sharing on mainframes; the virtual machine technology was developed by the industry, but has been finally superseded by modern operating systems. A discreet revival at the academic level took place in the 1980s, targeting time-sharing on massively parallel computers. The current trend has very strong industry support, and must be followed closely.

14.4. PROPOSED APPROACH

User requirements have to be translated into well-known computer engineering and computer science concepts, e.g. soft and hard real-time scheduling; connected mode vs. stateless or stateful services. With respect to Grid management, responsiveness translates into Quality of Service (QoS) requirements. Just as video rendering or music playing requires special scheduling on a personal computer, or video streaming requires network-differentiated services, Grid responsiveness requires the integration of a QoS constraint in scheduling policies and implementations. The implementation level is complex, but is nonetheless a continuation of what already exists. The policy level is more challenging, with a strong connection to usage policies. Examples from other contexts (processor scheduling, network routing) give indications that, in principle, QoS and best effort can indeed coexist. These contexts also provide a rich body of theory and experience.

The discussions held at the Linz e-IRG meeting brought out the following insights:

What type of benefits would a responsive Grid bring?

Detailed examples have been presented in the scenarios section above. In a nutshell, the consensus is that responsiveness might be critical for switching from community-based Grids to the full potential of the Grid paradigm, ‘the Grid for the masses’. However, there is a possible sociological obstacle on this path. There is little immediate incentive for the Grid decision bodies to focus effort on responsive Grids, and invest the time and manpower - which is significant - necessary to bridge the gap between current practice and this very motivating vision. Typically, the ‘early users’ step of the cycle will only marginally increase resource usage in comparison to traditional mammoth compute- or data-intensive applications, and thus will not show in statistics.

Do responsive Grids need a change of paradigm in scheduling, or are they amenable to the current Grid scheduling framework?

The consensus was in favor of an incremental approach. Production Grids, and especially EGEE, have proved flexible enough for the deployment of very interaction-demanding applications. Moreover, EGEE has demonstrated willingness to integrate responsiveness as a current priority.

From the body of experience, is it possible now to define a roadmap and a process for the integration of existing technologies (and particularly those developed by EU projects) into the production Grid middleware, enabling their widespread diffusion?

The answer to this question is more delicate. While the potential benefits are enormous, it is a fact that responsiveness has emerged only very recently as a well-defined topic. Various technical solutions do exist, and the investment in these developments has to be properly acknowledged and valorized.

Finally, indicators more refined than crude usage resource statistics, and more precise than user satisfaction, have to be defined in order to state explicitly the milestones of such a roadmap. ESFRI and e-IRG should be instrumental in working towards this goal.

14.5. RECOMMENDATIONS

1. The e-IRG recommends that awareness about the benefits of responsive Grids be raised, and that the above-mentioned issue of lack of immediate return on investment be alleviated by the creation of adequate incentives.
2. The e-IRG believes that coordination between the different EU projects addressing the issue of Grid responsiveness as a research and development task should be developed. The goal is to promote software interoperability, and factorization of developments.
3. A dialogue should be initiated between the Grid providers endorsing the objective of responsive Grids and the EU projects and actions in the areas of ubiquitous computing and ambient intelligence.
4. At the technical level, the virtualisation frameworks and tools and their impact on responsiveness should be surveyed. Because industry investment in this area is already considerable, close interaction with industry software providers is required.

14.6. TECHNICAL ANNEX

- Information on the EGEE SDJ working group can be found at http://egee-intranet.web.cern.ch/egee-intranet/NA1/TCG/wgs/sdj.htm.
- The first report of the SDJ working group is available at http://egee-intranet.web.cern.ch/egee-intranet/NA1/TCG/wgs/SDJ-WG-TEC-v1.1.pdf
- Information on the CrossGrid project can be found at: http://www.crossgrid.org
- A good summary on the work on interactivity is the chapter of the CrossGrid book that can be downloaded from: http://grid.ifca.unican.es/crossgrid/Archivos/InteractivityChapter_Final.pdf
- Information on the int.eu.grid project can be found at: http://grid.ifca.unican.es/int.eu.grid
- Information on the gridcc project can be found at: http://www.gridcc.org

15. LEGAL ISSUES IN E-INFRASTRUCTURES

15.1. INTRODUCTION

Ever since the 1980s, legal scholars and practitioners have been paying significant attention to the current and potential legal aspects of information technology. Starting with issues such as privacy and the legal protection of computer software, the debate and the statutory state of the art has grown to include topics such as electronic commerce, domain names, telecommunications, identity, security, liability, computer crime and legal protection of databases.

So far, e-infrastructures have not been discussed as a separate subject in the legal debate. This will have to change, given the growing size and societal importance of e-infrastructures.

This contribution is based on the discussion that took place in a session devoted to legal issues during the workshop of the e-Infrastructure Reflection Group, held in Amsterdam on May 13, 2005. The text reflects the outcome of the session, in which the main participants were technicians and policy makers. We both provide an overview of the legal issues as they are currently perceived and also draw attention to ways in which the legal discussion surrounding e-infrastructure might be advanced.

15.2. LEGAL ISSUES AS CURRENTLY PERCEIVED

E-infrastructures in their current form offer the possibility of exchanging and processing electronic data with unprecedented speed and capacity, involving a mixture of both pre-existing and new elements. On the one hand, we have well-known functionalities (electronic data processing, electronic communication), structures (electronic network) and parties involved (users, suppliers, intermediaries, governing bodies); on the other hand, there are new aspects such as the sharing of resources (with attendant questions of scale, consequences, preconditions) and the level of communication and processing capabilities.

If we take this division a starting point, an initial legal analysis of today's e-infrastructures will lead to:

1. the identification of legal issues that have been debated previously in relation to the Internet, mobile communications etc.;
2. the identification of legal issues that, although not entirely new, have a greater significance in an e-infrastructure environment than in a more 'traditional' computer network environment;
3. new issues, not encountered before in studying data processing and communication facilities from a legal perspective.

During the the Amsterdam workshop session, the following legal issues were perceived to be particularly relevant in relation to current e-infrastructures.

- *Intellectual Property Rights*: copyright and patent protection of software in particular might impede the development and adjustment of tools for operating and governing e-infrastructures;
- *Relations*: e-infrastructures are enabled by multi-party cooperation. This requires that adequate attention be paid to the structure, legal form and contents (roles and responsibilities) of the various relationships which can be distinguished;
- *Governance*: this includes issues such as responsibility, control, reporting and transparency;
- *Rules for sharing resources*, including Acceptable Use Policies;
- *Enforcement*: rules have to be implemented as well enforced in order to be effective;
- *Security*: which measures should and can be taken to ensure the integrity of the network and the data processing and who is responsible for taking these measures;
- *Authentication*: authenticating the provenance of instructions, statements, reports etc. is essential for safeguarding, amongst other things, the integrity of e-infrastructures;

- *Accounting*: facilitating the economic aspects of the use of e-infrastructures is key for its further development;
- *Competition Law Issues*: the (potential) economic significance of e-infrastructure requires careful consideration of competition law issues that might arise in the course of further development of e-infrastructures. The market for communication and processing capacity could be distorted by the capacity offered by e-infrastructures. This will be especially relevant when e-infrastructure facilities are made accessible to organisations with commercial interests, and state aid could be a relevant issue here. In addition, the rules for public procurement might be relevant when building e-infrastructures;
- *Responsibility and Liability*: virtual organizations play an important role in creating, operating and controlling e-infrastructures. This results in interesting questions regarding the extent to which they could be held responsible or liable for the way the e-infrastructures is being (mis-) used;
- *Criminal Law Issues*: processing of data in geographically distributed networks spread across various jurisdictions automatically leads to questions regarding criminal law responsibilities (e.g. applicable standards, which person or entity is responsible, jurisdiction). The proposals of the EU Ministers of Justice regarding data retention are also relevant here;
- *Data Protection*: e-infrastructures enable the exporting of data and extremely powerful data analysis, and either might endanger the privacy of individuals;
- *Applicable Law and Jurisdiction*: e-infrastructures are by nature cross-border and thereby multi-jurisdiction.

The participants in the legal session identified the following issues as being currently of specific importance: intellectual property, data protection, governance (including usage policies and the role of governments) and the validation of data.

It should of course be kept in mind that the concept of e-infrastructures is still developing and hence significant new insights might arise.

15.3. RECOMMENDATIONS

- The e-IRG proposes the organisation of a workshop aimed at identifying the legal issues specific to e-infrastructures. In this workshop legal experts in the field of information technology and law will analyse the specific features of e-infrastructures. Input to this workshop will be given by experts in the field of the technical and organisational aspects of e-infrastructures. The outcome of the workshop should support the decision-making process regarding the way that development of the legal component of e-infrastructures should progress.
- On the basis of the outcome of the Amsterdam workshop, the e-IRG recommends that an inventory is made of the legal problems which the scientific community currently feel impede the practical use of e-infrastructures.
- The e-IRG recommends that practical solutions for the legal problems currently being experienced be sought. These could range from (European) regulatory initiatives to standard-setting, developing certification schemes, codes of conduct, model contracts and other forms of self-regulation. In all cases, the issues should be addressed at the European level, and in some cases maybe even at a global level (treaties, standardisation initiatives, declarations etc.).
- The e-IRG believes it is time to start a more fundamental analysis of the legal issues involved in sharing computing resources and expanding the scope of this practice beyond the scientific community, both on the demand side (users) and the supply side (those making resources available).

- The e-IRG notes that legal issues should be an integral part of a multidisciplinary approach in the further development of e-infrastructures. Incorporation of legal aspects in the debate will be beneficial to the further development of e-infrastructures.

15.4. FOLLOW-UP

A workshop as recommended above is currently planned for the end of 2006 or the beginning of 2007. Other opportunities to foster the legal debate in relation to e-infrastructures will be monitored.

16. INTEGRATED DATA MANAGEMENT

16.1. INTRODUCTION

A task force to tackle the above area has been created during the UK presidency. A first draft paper has been presented during the UK e-IRG meeting in London and the task force continued its work during the Austrian presidency producing an updated version of the paper. During the Linz workshop a report was presented to the audience. However, there hasn't been possible to adapt the Task Force report to the template of the White paper and include it in this section for this version. This will be done later this year, during the Finnish presidency, and thus this sections acts as a place holder for the future work.

17. TERMINOLOGY

Acceptable Use Policy (AUP)

A set of rules restricting the ways in which a resource may be used. (From the wikipedia)

accounting

The tracking of resource usage by users and processes. Accounting is related to auditing, but is usually performed for allocation purposes.

allocation

The partitioning and provisioning of resources to entities (users and processes). For example, memory allocation refers to the allocation of memory between programs by the operating system. Grid systems allocate computing and storage resources to their users according to their policies, possibly taking into consideration accounting data.

ambient intelligence

The concept of ambient intelligence (or AmI) is a vision where humans are surrounded by computing and networking technology unobtrusively embedded in their surroundings. (From the wikipedia).

auditing

In computing, the processes (manual and automatic) that relate to the keeping of an inventory of the use of resources by users and computational processes. Auditing is related to accounting, but is usually performed for security purposes.

authentication

In computing, the process of verifying the identity of an entity (a person or a program).

authorisation

In computing, the process of selectively granting an entity access to a resource based on the entity's identity. The process is selective because it usually takes into account predefined authorisation policies.

CAOPS

The Certificate Authority Operations Working Group of the Global Grid Forum, which develops operational procedures and guidelines that facilitate the use of X.509 and other technologies for cross Grid Authentication.

co-allocation

The simultaneous allocation of a resource set. The resource set may comprise physically and geographically separated resources. (Adapted from the Globus toolkit).

eduGAIN

The GEANT Authorisation Infrastructure for the research and education community.

Eduroam

Eduroam stands for Education Roaming and is a RADIUS-based infrastructure that uses 802.1X security technology to allow inter-institutional roaming. See: http://www.eduroam.org

e-infrastructure

Electronic infrastructure, the distributed electronic resources necessary for Grid computing, storage, and networking.

e-science infrastructure

The electronic infrastructure that is required to enable e-science.

e-science

Computationally intensive science that is carried out in highly distributed network environments, or science that uses immense data sets that require Grid computing. The term was created by John Taylor, the Director General of the United Kingdom's Office of Science and Technology in 1999. (From the wikipedia).

European Research Area (ERA)

The research and innovation equivalent of the 'common market' for goods and services. It is a regrouping of all European Union supports for the better coordination of research activities and the convergence of research and innovation policies, at national and EU levels. (Adapted from the European Commission at http://ec.europa.eu/research/era)

externality

The effect of a transaction between two parties on a third party who has not played any role in the carrying out of that transaction. In economics therefore, an externality occurs when a decision causes costs or benefits to stakeholders other than the person making the decision, often from the use of common goods (for example, a decision which results in pollution of the atmosphere would involve an externality). In other words, the decision-maker does not bear all of the costs or reap all of the gains from his or her action. Externalities may be positive, when the result is beneficial (for instance, a house with a nice garden benefits others in the area) or negative (pollution as above).

Global Lambda Integrated Facility (GLIF)

An international virtual organization that promotes the paradigm of lambda networking. See: http://www.glif.is

Grid

Originally a "computational grid" was defined by Kesselman and Foster (1998) as a hardware and software infrastructure that provides dependable, consistent, pervasive, and inexpensive access to high-end computational capabilities, which translates into a networked hardware and software infrastructure that enables applying the resources of many (high end) computers to a single problem. As other types of Grids (data Grids, sensor Grids) have emerged the generic term Grid identifies a hardware and software infrastructure based on open standards that enables coordinated resource bundling and sharing within dynamic organisations consisting of individuals, institutions, and resources. The key concept is the ability to negotiate resource-sharing arrangements among a set of participating parties (providers and consumers) and then to use the resulting resource pool for some purpose.

Gridification

The process of porting an existing application to a Grid.

hard real-time

A system is said to be hard real-time if the correctness of an operation depends not only upon the logical correctness of the operation but also upon the time at which it is performed. In contrast to soft real-time systems, hard real-time systems will suffer critical failure if the time constraints are violated. (From the wikipedia).

Internet2

American consortium that develops and deploys advanced network applications and technologies for research and higher education. See: http://www.internet2.edu

Internet Engineering Task Force (IETF)

the standardisation body that produces and maintains the internet standards. IETF is part of the Internet Society. See: http://www.ietf.org

Internet Society (ISOC)

An international membership organisation that promote the open development, evolution, and use of the Internet. It is the organisational home of the Internet Engineering Task Force. See: http://isoc.org

interoperability

The ability of different systems to work together in order to accomplish a common task. With respect to software, the term usually refers to the ability of different programs to exchange data, use the same communication protocols, and employ the same file formats. (From the wikipedia).

Middleware Architecture Committee for Education (MACE)

A group of campus IT architects that provide technical advice and direction to the Internet2 community. See: http://middleware.internet2.edu/MACE

Metropolitan Area Network (MAN)

A class of network that typically covers an area of between 5 and 50 km diameter.

meta-scheduler

A meta-scheduler is responsible for balancing the workload between sites and data centres. Metascheduling is a technology that is responsible for managing jobs and application workflow, including submitting, scheduling, executing, monitoring, stopping, and retrieving results of computational jobs. (From gridwise tech, http://www.gridwisetech.com/metaschedulers).

middleware

Software that sits between the operating system and application programs and provides a set of common services to them. Although the term is thought to be new, it was actually used in the report of the 1968 NATO Software Engineering Conference that marked the birth of Software Engineering.

National Grid Initiative (NGI)

The national entity responsible for providing Grid access and services to the research and education community.

National Research and Education Network (NREN)

The national entity responsible for providing network access and services to the research and education community.

peer

An entity in a network that is constituted by similar entities, with no hierarchical relations between them.

peer-to-peer (P2P)

A peer-to-peer network is a network (often a virtual network on top of another, physical, network such as the internet) where the services and resources are provided directly by the users of the community. Peer-to-peer in its purest form does not rely on the services of centralised entities such as servers for its operation, but instead propagates service requests through an ad hoc topology and connects peers directly to each other.

Policy Management Authority (PMA)

A body to establish requirements and best practices for Grid identity providers to enable a common trust domain applicable to the authentication of end-entities in inter-organisational access to distributed resources. Coordinates a Public Key Infrastructure (PKI) for use with Grid authentication middleware. (Adapted from the EUgridPMA Charter).

porting

The activity of moving an application from one computer operating system environment (platform) to another.

production Grid

A Grid that is used for production work, in contrast with research Grids that are vehicles for research purposes.

public good

A good which, once produced, does not become scarce when it is used. One party's consumption of the public good does not reduce the availability of the public good to others. One example would be street-lighting.

Quality of Service (QoS)

The attainment of a service's goals, measured against pre-defined metrics. For instance, in networking it may refer to achieving a guaranteed throughput level between two nodes.

Remote Authentication Dial-In User Service

An Internet Engineering Task Force standard to authenticate users in a network

real-time scheduling

The scheduling of resources subject to a 'real-time constraint', that is, operational deadlines.

real-time

Hardware and software systems which are subject to a 'real-time constraint', that is, operational deadlines from event to system response. (From the wikipedia).

rivalrous

In economics, a good whose consumption by an agent reduces the amount of the good that is available for consumption by others.

schedulers

Software that plans the execution of processes on resources. Schedulers try to balance resource loads, take into account any policies that apply, and be fair to the requests they receive, subject to load and policy requirements.

Shibboleth

an Internet2 middleware project that provides an open, standards-based solution to the needs for organizations to exchange information about their users in a secure, and privacy-preserving manner.

soft real-time

A system is said to be soft real-time if the correctness of an operation depends not only upon the logical correctness of the operation but also upon the time at which it is performed. In contrast to hard real-time systems, soft real-time systems will not suffer critical failure if the time constraints are violated. (From the wikipedia).

supercomputer

A computer that is among the leading computers in the world in terms of processing capacity, at the time of its introduction. The term "Super Computing" was first used by New York World newspaper in 1929 to refer to a large custom built tabulators IBM had made for Columbia University.

t-infrastructure

Short for training infrastructure, i.e, specialised e-infrastructure for education including software, data and computer systems required to deliver courses and to support teachers and learners.

ubiquitous computing

The integration of computing elements into the environment, so that these become enmeshed in the fabric of everyday life instead of being distinct objects.

Virtual Organisation (VO)

A group of individuals or institutions collaborating together on a specific issue and facilitating this by sharing information and other resources (for instance, through Grid technology).

virtualisation

The process of presenting a subset of computing resources so that they can be accessed in ways that give benefits over the original configuration. This new virtual view of the resources is not restricted by the geographic location, the physical configuration or the implementation of underlying resources. Commonly virtualized resources include computing power and data storage. (Adapted from the wikipedia)

visualisation

The representation of data (whether quantitative, qualitative or metadata) or conceptual models through some means of graphical rendering in order to make it easier to analyse.

18. ABBREVIATIONS

AA
: Authentication & Authorization

AAA
: Authentication, Authorization, and Accounting

AAI
: Authentication Authorization Infrastructures

AC
: Attribute Certificate

ACL
: Access control list

APGridPMA
: Asia Pacific Policy Management Authority for Grid Authentication, see: http://www.apgridpma.org

AUP
: Acceptable Use Policy

CA
: Certificate Authority

e-IRG
: e-Infrastuctures Reflection Group, see: http://www.e-irg.org

ERA
: the European Research Area

E2E
: end-to-end

EGEE
: Enabling Grids for E-sciencE, see: http://www.eu-egee.org

ESFRI
: European strategy Forum on Research Infrastructures, see: http://cordis.europa.eu/esfri

EUGridPMA
: European Policy Management Authority for Grid Authentication, see: http://www.eugridpma.org

GGF
: Global Grid Forum, see: http://www.ggf.org

GLIF
: Global Lambda Integrated Facility

GN2
: GÉANT2, see: http://www.geant2.net

GOC
: Grid Operation Centre

HPC
: High-Performance Computing & Networking

HPC

High-Performance Computing

ICT

Information and Communication Technology

IETF

Internet Engineering Task Force

IGTF

International Grid Trust Federation, see: http://www.gridpma.org

IP

Internet Protocol

IPR

Intellectual Property Rights

IPv6

Internet Protocol version 6

ISOC

the Internet Society

IST

Information Society Technologies

LCG

LHC Computing Grid

LHC

Large Hadron Collider, a particle accelerator located at CERN in Switzerland/France.

MACE

Middleware Architecture Committee for Education

MAN

Metropolitan Area Network

MoU

Memorandum of Understanding

NGI

National Grid Initiative

NOC

Network Operation Centre

NREN

National Reseach and Education Network

NREN-PC

Policy Committee within GEANT2 with appointed representatives from each partner in the project.

P2P

peer-to-peer

PMA

Policy Management Authority

RADIUS

Remote Authentication Dial-In User Service

SME

Small and Medium Enterprises

SLA

Service Level Agreement

TAGPMA:The Americas Grid Policy Management Authority, see

http://www.tagpma.org

TEN

Trans-European Network

TACAR

TERENA Academic CA Repository

TERENA

Trans-European Research and Education Networking Association is an association of organisations that are involved with the provision and use of computer network infrastructure and services for research and education in Europe., see: http://www.terena.nl

TF-EMC2

Task Force on European Middleware Coordination and Collaboration (within TERENA).

VO

Virtual Organization

VPN

Virtual Private Network

19. ANNEXES

19.1. WHITE PAPER CONTRIBUTORS

White Paper Chief Editor: Fotis Karayannis fkara@grnet.gr

White Paper Co-editor: Panagiotis Louridas

Editorial Support: Joanne Lawson

Section Editors (wp-editors@e-irg.org):

1. Preface: Dieter Kranzlmüller
2. Introduction: Panagiotis Louridas
3. Executive Summary: Fotis Karayannis, Panagiotis Louridas
4. Recommendations: Section editors
5. Authentication and Authorisation Policies: Diego Lopez, David Groep
6. Research Networking Policy Challenges for e-Infrastructures: Afrodite Sevasti, Lajos Balint, Vasilis Maglaris
7. Incentives for Resource Provision in Scientific Grids: Panayotis Antoniadis, Costas Courcoubetis
8. Usage Policies: David Kelsey
9. Training and Education: Malcolm Atkinson, David Ferguson, Wolfgang Gentzsch, Brendan Hamill
10. Grid Economy – Allocation and Accounting:Anders Ynnerman, David Kelsey
11. Middleware Interoperability and Repositories: Steven Newhouse
12. User Support: Frank Harris, Torsten Antoni, Harald Kornmayer
13. Supercomputing and Petascale supercomputing Harmonisation Policies:Sergi Girona, Victor Alessandrini
14. Responsive Grids – Short Jobs and Policy Issues: Cecile Germain-Renaud, Massimo Lamana
15. Legal Issues: Kees Stuurman
16. Integrated Data Management: Alan Foster

17-18: Terminology - Abbreviations: Fotis Karayannis, Joanne Lawson, Panagiotis Louridas, Michiel Leenaars

19. Annexes: Fotis Karayannis, Panagiotis Louridas

Written comments provided by Matti Heikkurinen, Bob Jones, Fiorenzo Scaroni, Torbjörn Wiberg.

19.2. WHITE PAPER CONCEPT TEMPLATE

This document presents the layout of the contributions for the e-IRG workshop that will serve as the basis for the Austrian e-IRG White paper.

Policy Area: A short paragraph describing the policy topic under discussion.

Goal: A short paragraph describing what the White paper chapter is about and how it relates to the overall theme of e-Infrastructures policies.

Context: Current practice – achievements and limitations: A short section describing recent advances and limitations in the area, i.e., policy state of the art; more technical material should be placed in an appendix. It should highlight the main issues within the current situation that need to be tackled. The section should end with a list of challenges in the form of questions to be explored.

Questions to be answered during the workshop:

- Q1
- Q2
- …

Proposed approach: If there is a foreseen solution or proposal under consideration, this should be presented in this section. Otherwise the workshop should function as a venue for ideas and discussion.

Recommendation: This section should be prepared following the results of the workshop and as part of the White Paper editing procedure, and should be ready for the June 2006 e-IRG meeting. If a mature proposition is available at an earlier stage, this should be put forward.