# eInfrastructure Reflection Group

Setting the scene through a first White Paper

Version 2.3

4 December 2003

## Table of Contents



# 0

# Background – The New Environment of eInfrastructures

The explosive growth of technologies associa ted with computing and electronic communication is providing unprecedented opportunity for growth and change in society and has the potential to drive economies based on information and knowledge. Applications currently being developed in this dynamic env ironment require ubiquitous distributed electronic infrastructures, dubbed eInfrastructures, created from the integration of existing and developing research networks, large - scale computing fabrics, and nascent grid middleware environments. The exploitati on of the enormous potential computing, storage and other electronic resources within the various user communities linked by broadband optical networks is of key importance for the European Union and international scientific and economic communities. Init iatives such as the EU eInfrastructures initiative will make this vision a reality.

To date, the World Wide Web has provided transparent access to information for millions of Internet users. The new electronic infrastructures are intended to extend this to provide rapid, secure, and transparent access to distributed computing resources and services. This "World Wide Grid" of resources will form the basis of the Information and Knowledge Society, and will be built upon the software and hardware necessary t o establish virtual collaborative environments, tools for education and research, planning and simulation tools for complex problem solving, economic modelling analysis tools, virtual environments for medical treatment, storage and analysis of high - resolution digital data, pictures, and video and for providing access to massive scientific databases for disciplines from bio-informatics and bio-chemistry to meteorology, physics, and astronomy.

Pioneering work in these areas has been done by the academic rese arch and scientific communities. These electronic science (e -Science) applications are building the frameworks and creating the necessary impetus for the growth of the required architectures and standards. At the same time, these Grid technologies are being adopted by the wider community of the Information Society, with applications such as e -Government: civilian transactions with administrations and governments, e - Business: providing tools and services for business, and areas such as financial modelling, data storage and analysis for medical and pharmaceutical sciences, entertainment and advertising, and the simulation of complex technological systems.

Today, electronic infrastructures are implemented through grids of computing and storage resources conn ected through electronic networks of local, national, and international scales. The field is now learning how to transform these research environments into production -quality infrastructures capable of supporting these communities.

These new technologies provide unprecedented opportunities for novel means of education, economics, collaboration, and scientific endeavour amongst others. However, they also bring new issues regarding policies that must be understood in this new environment in order to exploi t their full potential. Such issues as models for acceptable resource sharing and accounting of the associated cost, entitlement of communities or individuals to access and use resources, responsibility, privacy, to name but a few, cross traditional natio nal, economic, and political boundaries. This paper is intended as a first look at such issues, providing an overview of the current state of the art, bringing together experiences and knowledge gained by the current generations of grid projects.

## *0.1 The EU eInfrastructures initiative*

For the development and support of the eInfrastructures environment a series of workshops has been launched by the European Union under the aegis of the European Union Presidencies - http://www.einfrastructures.org. The "eInfrastructures" paradigm will reach its broadest scope and cross-border relevance, with administrative and policy decision mechanisms that will satisfy the

diverse end-user communities' requirements of performance, service transparency and security, while achieving scale economies in providing ever-growing resources at attractive cost.

### 0.1.1 Athens, 12 June 2003 - Launching the eIRG

On the 12th of June 2003 under the auspices of the Greek presidency of the EU, the 1st workshop was held in Athens, organised by the General Secretariat for Research & Technology (GSRT), the European Commission and the Greek Research and Technology Network (GRNET) in collaboration with the Greek National Documentation Centre (EKT). The workshop, entitled Towards integrated Networking and Grids infrastructures for eScience and beyond -The EU eInfrastructures Initiative, aimed at discussing the creation of the necessary administrative and policy decision mechanisms for the successful deployment of "eInfrastructures" within the extended European Research Area. Among the key recommendations of the workshop was the establishment of an **eInfrastructure Reflection Group (eIRG)** with a membership "built from national representatives". The eIRG "should consider and communicate clear messages on policy issues to both European Commission and existing infrastructure projects".

### 0.1.2 eInfrastructure Projects – Double Role of Generating and Adopting Policies

In this respect it is important to note that two significantly sized initiatives will be launched in the context of FP6 eInfrastructures, which are expected to help structuring the grid infrastructures in Europe and building upon the already established GEANT infrastructure. These projects – EGEE / SEEGRID and DEISA – will be major actors in this context as they both generate policies and are major players in adopting and generalising them.

- **EGEE- Enabling Grids for E-Science in Europe,** which is led by the European Centre for Nuclear Research (CERN), aims to build the largest international grid infrastructure to date, operating in more than 70 institutions throughout Europe, providing 24-hour grid service and a computing capacity comparable to 20,000 of today's most powerful personal computers. In EGEE a specific networking activity has been proposed to assist the work of the eIRG and a selected number of partners have been committed to promote it.
- **DEISA**, which will build and operate a distributed terascale supercomputing facility, whose integrated power will be close to 30 teraflops in 2004. The principle objective of this project is to advance computational science in leading scientific and industrial disciplines, by deploying an innovative GRID-empowered infrastructure to enhance and reinforce High Performance Computing in Europe. The proposed infrastructure is based on the tight coupling – using dedicated network interconnects – of six homogenous national supercomputers, to provide a distributed platform and is based on an innovative operational model, capable of providing substantial European added value to the existing national infrastructures. The distributed multi-cluster platform is in turn integrated into a larger heterogeneous Grid.

### 0.1.3 Rome, 9 December 2003 – Consolidating the eIRG and Defining the Scope

On the 9th December 2003, the 2nd eInfrastructure Open Workshop entitled: "eInfrastructures (Internet and Grids) - The New Foundation for Knowledge-Based Societies" will take place organized by the Italian Ministry for Education, University and Research (MIUR) under the aegis of the Italian Presidency of the European Union with the High Patronage of the President of the Italian Republic, Carlo Azeglio Ciampi, and in co-operation with the European Commission. The key objectives of this workshop include:

- Deciding on concrete next steps and actions required for the establishment of a policy - and administrative-level framework in Europe and beyond, across technological, administrative and national eInfrastructure domains,
- Facilitating the first meeting of the eInfrastructure Reflection Group.

One of the primary goals of the meeting, at which key players in the construction of the EU eInfrastructure will be present, is to review the perspectives and the technical and political issues related to the usage of the eInfrastructure for Science and Society at the national, European and international level.

## *0.2 The eIRG Meeting and an Initial Draft White Paper*

The first meeting of the eIRG will take place at the INFN Headquarters on December 10th 2003, the day after the Open Workshop. The main objective of the meeting will be to decide on the Terms of Reference of the Group and an operational structure and work-plan.

An initial White Paper is being drafted as a first input to articulate the discussion and is expected to become a live document to continuously support and reflect on the work of the eIRG. The editorial board of this White Paper currently comprises the following members:

- Victor Alessandrini (CNRS-DEISA)
- Kyriakos Baxevanidis (EU)
- Ian Bird (CERN-LCG)
- Alan Blatecky (SDSC)
- Brian Coghlan (TCD-Ireland)
- Fabrizio Gagliardi (CERN)
- Francois Grey (CERN)
- Fotis Karayannis (GRNET) – editor
- Peter Kaufmann (DFN)
- Dave Kelsey (UK GridPP)
- Mirco Mazzucato (INFN)
- Jesus Marco, (UNICAN-CSIC)
- Federico Ruggieri (INFN)
- Matti Veikko Johan Heikkurinen (CERN)

*This editorial board will in future be supported by the EGEE project (especially in the context of the Networking Activity 5, coordinated by Fabrizio Gagliardi from CERN). DEISA, SEEGRID & other FP6 projects, as well as GEANT are also expected to actively participate.*

The White Paper will summarize current achievements concerning general policies in use to address resource access and sharing at the pan-European and International level across different administrative and national domains (a useful input is the preliminary work from the particle physics community in the LCG project and other experiences), focus on the major issues that need to be addressed, such as accounting, a cost model etc., and identify a roadmap for future development of a political and administrative framework in Europe and Internationally in order to allow a real effective exploitation of the eInfrastructures.

# 1 The Scope of the White Paper

## 1.1 Introduction

The White Paper summarizes the current status of the general policies in use concerning resource access and sharing at pan-European and International level across different administrative and national domains (starting from what is being implemented within major Grid projects, like the EU Datagrid project and the major Particle Physics Grid project – LCG, other national Grid Initiatives such as the UK e-Science programme, the Italian INFN grid, NorduGrid, and what has been achieved in major Supercomputing centres and test-beds, e.g. the Teragrid in the US). It will focus on the major issues that need to be addressed in this domain: authentication, authorisation, accounting, cost and business models etc. It will identify a roadmap for the future development of a political and administrative framework in Europe and internationally to allow a rapid deployment and effective exploitation of the eInfrastructures.

## 1.2 eIRG Terms of Reference

The main objective of the eIRG is 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.

The eIRG will consist of high-level national experts and national representatives that will have the mandate from their national public research authorities to participate in the above group and represent their country's views on the above topic.

The work of the group will be assisted by relevant initiatives at the technical level (e.g. FP6-funded efforts) that will conduct the necessary technical work to support the effective performance of eIRG tasks (e.g. creation of a registry of access and use policies of electronic resources in Europe, preparation of white papers on the harmonization of resources use policies etc).

## 1.3 Methodology

The proposed methodology constitutes the basis of the proposed document. The policy framework will be the outcome of two main streams. The first stream will be the usual "use-model, requirements-capture, use-case-specification, architecture" chain, where the entities and their roles are first identified, their requirements are captured and analysed along with use case specifications resulting in the policy framework layout. The second stream is the capture and analysis of current practises and achievements in resource access and sharing. The combined analysis of the two streams will provide the policy framework and will include a series of policies.

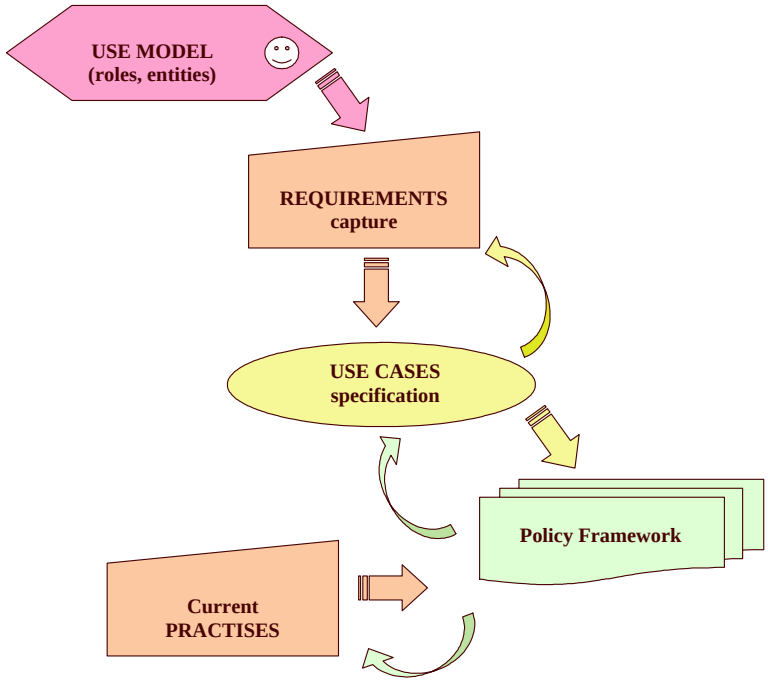
USE MODEL
(roles, entities)
REQUIREMENTS
capture
USE CASES
specification
Policy Framework
Current
PRACTISES

# 2 Use Model

## 2.1 Introduction

This section provides a draft use model (or business model) including all entities (or actors) of the eInfrastructure (e.g. the users, Virtual Organisations (VO), middleware providers and operators, resource providers and operators, network providers and operators etc.). At a later stage their interactions will be depicted in a corresponding diagram. The administrative domains should also be visible. In this section the roles of the different entities and their responsibilities will be highlighted. Finally some basic use cases will be elaborated (to be refined after a detailed requirements capture and analysis following a questionnaire). At this first attempt significant input has been received from the corresponding work of the LCG Security Group on *Security and Availability Policies for LCG.* Other projects or national initiatives are welcome to participate.

## 2.2 Entities – Roles and Responsibilities

### 2.2.1 Grid Architecture

Different frameworks for a "production-quality" operation of Grid-aware distributed computing systems have been discussed in many projects and in dedicated sessions at workshops. These include the experience in current large test-beds, like EDG, CrossGrid, LCG; the Grid-Start inventory document; the London workshop on VO May'03, and the Brussels concertation meeting in June'03; and the Production Grid Management, Grid Economic Services Architecture, and SAAA GGF working groups; and the chapters on "Grid resource allocation and control using computational economies" in the book "Grid Computing".

Grids are defined as frameworks enabling coordinated resource sharing and problem solving in dynamic, multi-institutional virtual organizations. The framework is expressed with a layered Grid architecture, as is the case in the Lightreading report on the "Architecture of the Grid" shown in the diagram below. The specific architecture has been preferred since the different layers correspond to discrete physical components and can be easily linked with appropriate entities and roles. The basic layers of a Grid Infrastructure are the following from bottom to top:

A. The **network layer**, providing the physical interconnectivity for the Grid components using the respective networking equipment (routers, switches, etc.)

B. The **resources layer**, providing the actual Grid fabric resources excluding network resources belonging to the network layer, such as computational, storage or other (e.g. sensors, telescopes, etc). The resources are part of a **resource center.**

C. The **middleware layer**, providing all the protocols and components enabling the sharing of resources. The basic middleware components facilitate among others information services, resource allocation and scheduling, security emphasizing on authentication and authorization, monitoring and discovery services etc. The middleware layer encompasses the "connectivity" -, "resource" and "collective" layers according to most other layered architectures as the in "Anatomy of the Grid".

D. The **application** & Serviceware Layer, providing the actual **applications** of different scientific or business fields that will use the Grid, along with supporting portals and development toolkits.

The architecture should be enhanced with the distinction of **collective** vs. **non-collective** middleware layers, since this is closely related with the likelihood of different administrative domains providing these two services. In other words, non-collective middleware means that a

single entity can publish services of its local resources and even charge by itself; while collective middleware is about publishing services and managing multiple resources across domains.

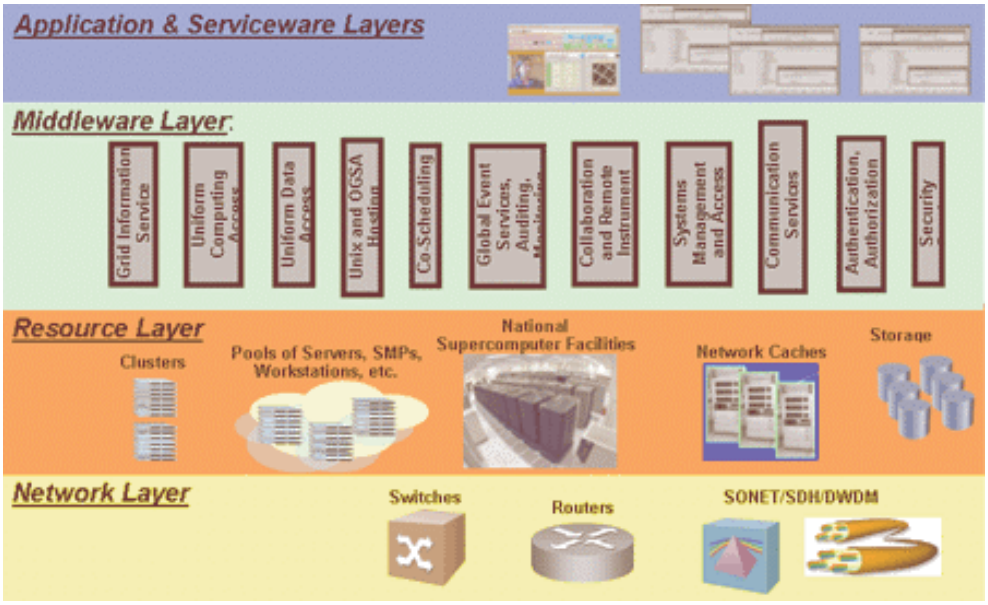


**Main components of a Grid** (Source: Lightreading)

According to the definition of the eInfrastructure, the 3 bottom layers are those constituting the **eInfrastructure. The users or consumers interact with the applications that sit on top this.**

### 2.2.2 Entities-Roles

The distinction between "users" or "consumers" and "providers" is a basic principle for discussing the sharing and access to resources . This is important because the "computational economy" is usually blurred - especially in an e -Science environment where part s of the grid computing resources are directly operated by the user communities themselves. Even when this is not the case, these resources are usually funded at national or institutional level, with the primary objective of satisfying the needs of their national or institutional community. In the academic framework, there is also a long tradition on sharing resources between different projects.

Administrative domains (national, local, etc) must also be taken into account. The current experience on the operation of network infrastructure s can provide an indication of the path to follow.

Keeping in mind the above architecture and basic principles, we can identify the following roles:

- **Application Layer**
  - **Users** (or end-users or consumers**),** interacting with the application and service-ware layer, running their applications on the Grid and actually utilising the eInfrastructure. The applications are usually related to electronic Science (e - Science), but there are some first examples of electronic business or government (e - Business, e-Government), where the use of a Grid will be expanded at a later stage. In our research and education environment the purpose of eInfrastructures is for e - Science. Users should be authorised as registered members of VO s and are able to obtain suitable authentication credentials containing their identity which have been

signed, directly or indirectly, by one of the Certification Authorities recognised by the eInfrastructure.

- **Virtual Organis ations (VOs) ,** which is a (sub -)group or association of users collaborating in a common experiment , project or other a joint venture. The Virtual Organisation is one of the most important roles in the eInfrastructure and adds great complexity, since VO s are formed as a selection of users belonging to different administrative domains . VO users are legally bound to the institutions where they work (in an e -Science framework usually working in different projects through **collaborations**, with a financial framework defined through a MoU or Annex ). So, legal responsibilities need the use of individual authentication, with certificates explicitly including the legal Institution, while resource usage is contemplated at the Collaboration/Project level.

- **Middleware Layer**
  - **Middleware providers (or developers) ,** implementing, testing, supplying and maintaining bundled releases of the necessary software. Besides the functional characteristics of the software, special attention has to be paid to non -technical characteristics such as robustness and reliability as well as to the security aspects . Currently there are multiple middleware packages as well as different versions of the same package. Thus, an important aspect of interoperation among eInfrastructures is the use of standard interfaces and/or common middleware package s. Open Middleware Institutes Initiatives that serve interoperability purposes are of key importance for eInfrastructures. As mentioned it is important to distinguish between collective and. non -collective providers, since this might involve different administrative entities . However, t his is much more important for the operation al aspects.
    - **Non-collective** middleware providers: This includes the publication of the available resources, the reservation and access by authorized users, as well as related accounting services.
    - **Collective** middleware providers, including among others meta-directory services, co-allocation/co-reservation/brokering services, replica location services along with monitoring and diagnostics services.
  - **Middleware operators ,** responsible for operating the Grid middleware in an eInfrastructure as part of their **Grid Operation Centers** (GOCs) and guaranteeing an agreed quality level . It is obvious that m ultiple GOCs in a Grid are mandatory for scalability and redundancy reasons; however different GOCs can focus on the provision of different services. From a simple economy perspective, GOCs could play some of the roles of a distributed computational market operator. From an administrative point of view, they can be assigned to a single legal entity, or can be organized in a distributed way with different partners offering different services, under a common project or a signed MoU.

    Inside EGEE there is a hierarchy of GOCs. The non-collective services are offered by the Resource Centres, while the Regional Operation Centres provide a collective service to support middleware installation and user support for a set of Resource Centres that are geographically close. Note that this can n ot be defined as a "traditional grid service" that can be invoked inside an application. Core Infrastructure Centres are centrally located and provide most of the collective services, including the Information and Replica Index, or Resource Brokering, and also collect the accounting information from the set of Resource Centres they serve. At a higher level, the Operation Management Centres (OMC) (GOC, as designed in the LCG project) provide services like the coordination of Grid operation, definition

of Service Level parameters, coordination of security activities and also monitoring of service performance levels.

- **Resource Layer**
  - **Resource providers (or resource centres)** supplying the Grid fabric resources. The grid resources are either part of an organised resource centre (RC) (with computer clusters or supercomputers), or they are individual resources part of a desktop Grid. The resource centres will allow access to all or part of their CPU/storage resources (as computing centres), through grid-specific network transactions. Resource centres can range from small (departmental) to large multi-function facilities. From an administrative point of view, they have a well-defined legal status and identity. From an economy perspective they can also be seen as a basic entity on the "production" side.

  - **Resources (or fabric) operators or administrators,** responsible for operating the Grid fabric in the organised resource centres. In case of desktop resources, the operation of the resources is outsourced to the resource centre operators or to the administrator of the department where the resources are located. In case of organised resource centres a designated resource centre manager is the interface with the rest of the world.

- **Network Layer**
  - **Network providers** supply the underlying networking connectivity. The adopted model in the research networking community is hierarchical. The resource centres are usually connected to a campus LAN (inside a University or Research organisation), the campus LAN is interconnected with the National Research and Education Network (NREN) and the NREN is connected to the pan-European Research and Education network (GEANT). End-to-end connectivity relies upon multiple administrative domains i.e. campus-NREN-GEANT-NREN-campus and in some cases there is also an additional regional network, either inside a province in a country or among multiple countries. The network providers rely mostly on leased capacity from national or international carriers but in some cases they own their networks (having acquired Indefeasible Rights of Use in optical fibres from third party carriers or building their networks from scratch. Service level agreements exist between the research networks and the carriers in terms of service availability, Mean Time Between Failure (MTBF), Mean Time To Restore (MTTR) etc. SLAs between GEANT and NRENs do not exist yet and will be studied during the GN2 lifetime.
  - **Network operators and administrators**, responsible for operating the networking part of the eInfrastructure as part of their Network Operation Centers (NOCs). Due to the multi-administrative domain environment, it is obvious that there are multiple NOCs, each responsible for its own domain. In other words, GEANT, Regional networks, NRENs or Campus networks operate individual NOCs, which cooperate with one another under agreed procedures. Service level agreements are used to guarantee an agreed quality level. Network operators, in cooperation with the corresponding administrative structures of each NREN and of GEANT, are responsible for safeguarding their existing Acceptable Use Policies. Note that GEANT is content with the national Acceptable Usage Policies (AUPs) of each NREN and does not possess its own AUP. However, in order for an NREN to connect to the GEANT network a list of high-level requirements must be met. (http://archive.dante.net/geant/connect.html)

### 2.2.3 Responsibilities

Based on the above entities and roles we can identify the following responsibilities per category (based upon input from the LCG Security Group):

- Users:
  - Safeguard Credentials and Private Keys: Users must ensure others cannot use their credentials to masquerade as them or usurp their access rights. The holder of a private key will be held responsible for all actions, whether carried out by the holder personally or not, carried out using creden tials generated from that key. No intentional sharing of credentials is permitted.
  - Observe Access Controls: Users must be aware that their jobs will oft en be running on equipment and using resources owned by others. They must observe any restrictions on access to resources that they encounter and must not attempt to circumvent such restrictions.
  - Observe Limitations on Use: Resources may be used only for legitimate professional purposes connected to the purpose of the eInfrastructure. Personal use of any nature is expressly forbidden.
  - Applications: Applications software written or selected by Users for execution using the eInfrastructure Resources must be directed exclusively to the legitimate purposes of the latter. Such software must respect the autonomy and privacy of the host sites on whose Resources it may run.
  - Respect for Others: Users must be aware that their work may be utilising shared resources and may seriously affect the work of others. They must show responsibility, consideration and respect towards other users in the demands they place on the eInfrastructure.

- Virtual Organisation:
  - User Registration: VOs are required to set up and operate a set of Registration Authorities and associated procedures for approving requests for joining the VO in accordance with the directions laid down by the eIRG. Approval must be restricted to individuals who are recognised as having legitimate rights to memb ership. V Os are subsequently required to maintain the accuracy of the information held and published about their members, and to promptly remove membership from individuals who lose their right to membership.
  - Controlling Access to Resources: Some resource s will be restricted to all members of certain V Os or to certain individuals within V Os. V Os will provide access to information about their members as necessary to enable such controls to be implemented and maintained accurately.
  - Applying Sanctions to Use rs: VOs are responsible for investigating reports of users failing to comply with the provisions of this Policy, and for taking appropriate action to ensure compliance in the future. This action may include the notification and involvement of the User's h ome institute. The ultimate sanction to be exercised at the discretion of the VO is the removal of membership, and hence the withdrawal of rights of access to the eInfrastructure resources.

- Middleware providers
  - Facilitating Security Controls: The softwar e should implement appropriate security techniques to control access to resources of all types.
  - Maintaining the Integrity of Services: Before distributing replacements, upgrades or patches to existing software, developers must ensure that adequate testing is carried out to ensure the functionality and reliability of existing Services will not be

jeopardised. When carrying out tests, developers will follow current best practice. This requirement may be relaxed if it is imperative that a security -related patch be distributed urgently.

- Middleware operators:
  - Contact details: The GOC is responsible for maintaining contact details of security personnel at each participating resource centre and for facilitating eInfrastructure-related intercommunications between them.
  - Monitoring SLAs: The GOC is responsible for monitoring the operational performance of the eInfrastructure services and for publishing details of its findings for comparison with the published SLAs of those services.
  - Security Expertise: The GOC together with the eIRG is responsible for establishing and maintaining expertise in eInfrastructure-related aspects of security in order to provide detailed advice and guidance to the community on avoiding and responding to internet security incidents.

- Resource providers
  - Quality Services: RCs accept the responsibility for providing quality services to their users.
  - Risk Assessment: RCs providing resources to the Grid acknowledge the risk of intrusions and host compromises and are responsible for assessing and minimising the risks. RCs should take the necessary measures to safeguard their resources.
  - Cooperation: In case of security incidents, RCs accept the duty to cooperate with the other structures in order to investigate and resolve the incidents taking the appropriate actions and sanctions.

- Resource operators
  - Site Policy: Resource operators must ensure their implementations of services comply with both their RC policies and this Policy.
  - Notifying Site Personnel: Resource Administrators are responsible for ensuring that all appropriate personnel concerned with security or system management on their site are notified of and accept the requirements of this Policy before implementing any services.
  - Resource Administration: The Resource Administrators are responsible for the installation and maintenance of Resources assigned to them, and subsequently for the quality of the operational service provided by those Resources. This quality will be defined by the Service Level Agreement (SLA) for each Resource as published by the Administrator of that Resource.
  - Service Level Agreement: The Administrator of each service instance must maintain an assessment of the risks inherent in their particular Service design or resulting from local services or operational practice which might affect that Service's Availability, Reliability or Performance, and publish the expected values of these service parameters in accordance with the GOC Procedures for Resource Administrators.

## 2.3 Use Model

*TBD*

## *2.4 Use Cases*

We categorize Use Models into several groups, starting from the large-scale, homogeneous ones, and proceeding through intermediate steps down to an individual user level. As a general remark, we note that, for large scale corporations, Grid usage patterns are fairly predictable, whereas approaching the scale of SMEs and individual users, Grid usage is anticipated to be more complex. Unpredictable behaviour of the smaller scale users poses additional requirements that need to be satisfied before economies of scale – the typical benefits of Grids for international corporations – can be achieved. The specific challenges involved for small scale users can be technical in their nature, but also they can be related to the issues of trust between individuals and organizations.

### 2.3.1 International corporate environments

From the cost-benefit analysis point of view, an international corporation provides the easiest environment for deploying Grid technology. One indication of this is that several major deals have been made over the last couple of years involving provision of Grid infrastructure by companies like IBM, to major corporations such as Royal Dutch Shell, which these corporations use for their data processing. Note that, according to Ian Foster, such corporate environments are not considered Grids, because all branches belong to the same administrative domain and the sharing problem is less complex. However, this will appear in new projects where collaboration among different corporate administrative domains entities will take place.

The current basic Grid solutions enable global companies to optimize the use of their resources by pooling them logically into a single resource pool. The benefits of this approach are partly related to the ability to defer hardware acquisitions through more efficient use of resources that the company has. Especially in the case of companies spanning several continents, the ability to ship computing to different time zones enables the company to balance the daily load cycle (get extra capacity from areas where it is night time). However, there are several additional benefits from this pooling of resources that could be an even bigger factor in promoting Grid technology in these kinds of environments:

- By making resources location independent, it is easier to manage maintenance breaks in individual computing centres. Grid technology makes it possible to add and remove resources dynamically from the network, thus as long as the total capacity of the computing infrastructure is sufficient, removing one centre from the resource pool does not prevent the rest of the company from doing work.
- The location independence will also make it possible to concentrate new hardware acquisitions to a smaller number of sites, with obvious economies of scale benefits. If instead of 10 smaller bids spread out throughout the world the company can ask for a single one, it is likely that the competition for this larger bid will reduce the amount of money the company pays per single computing unit.
- The total amount of manpower needed for monitoring the infrastructure can be reduced. If some of the routine maintenance tasks can be made location independent through various fabric management techniques, it is not necessary to have round-the-clock personnel present on all the sites. Using the differences in time zones to reduce the amount of work performed at night time is also an attractive way to reduce labour costs.

Achieving these benefits is dependent on several factors, which define the applicability of this model:

- The company must be involved in a business where it is either performing large amounts of relatively independent computationally intensive tasks – or where it needs to process large amounts of data. If the main part of company's computing is related to normal office work (word processing etc), the current Grid technology is not yet able to provide cost-

effective solutions, even in cases where the corporate entity is larger than any of the scientific collaborations developing the Grid.
• The network connectivity on the sites and between sites must be of sufficient capacity, and the savings from using the Grid must cover the costs of extra data traffic between the sites. On the other hand, if the company's main business is selling network connectivity and computing capacity as a service to other companies, the Grid holds considerable promise as a way to minimize the amount of extra hardware the company needs to purchase as a safeguard against local peaks in the resource usage.

A potential limit to a widespread use of Grid technology in this environment is the often high sensitivity of corporate data and the reticence to moving them out of the corporate computer glass room. For this the Grid will have to improve the level of security and offer sophisticated levels of access control.

### 2.3.2 Established International Scientific Collaborations

Certain international scientific collaborations rival or exceed international companies in terms of number of participants, amount of data to be processed, complexity of the organization and number of sites. In the case of high-energy physics, genome research and satellite imagery, the amount of data that needs to be processed is orders of magnitude larger than in typical corporate environments. While the storage space itself can't be shared by more than one user at the time, pooling of resources makes it possible to concentrate labour-intensive activities like backup copies to a subset of sites participating in the collaboration. Also writing the data on several geographically separated locations offers advantages both in terms of the reliability of the infrastructure and of the optimization of the use of interlinked commodities of CPU time, network bandwidth and mass storage.

Achieving these benefits requires developing an additional layer of virtualization compared with the typical corporate environment. In order to group resources from several distinct organizations into a single pool, the user identities in each of the participating organizations need to be mapped to a global identity (most often a Grid certificate). Similarly, collaborations need to decide what kind of operations users should be permitted to perform. In the case of big science collaborations, the extra cost and effort related to the maintenance of this mapping of local and global identities is seen as acceptable, since there is no other way to provide sufficient infrastructure for the scientific tasks the communities want to perform.

In the two above cases a major issue to be addressed is the interconnectivity between the local network fabrics and the WAN backbone. Together with the issue of last mile connectivity this is often the bottleneck to achieve a widespread and geographically distributed end user community.

### 2.3.3 Academic computing service providers

Academic service providers maintain networks that connect their member institutes and connect this network to the other networks. Academic service providers tend to have a somewhat different approach from companies in calculating the cost of their operations and prices charged from their member organizations. The cost sharing model might be based on the GNP of nations participating on the network (an example is Nordunet) or may apply a more complex costing model that takes into account the connection speed of the member country and the underlying international connection costs of the country in question.

Grid applications change some of these assumptions, since the load will likely be concentrated on the links between major computing centres. Furthermore, large amounts of the traf fic on these links can be generated relatively independently from the speed of the connection to the country of origin through e.g. copying of input and output files to permanent storage.

One of the open issues is the effect of the differences in the cost structure of individual institutes participating in Grid research. The cost sharing arrangements inside each country tend to be different from each other and depended partially on relatively fixed parameters (e.g. size of the organization) and the amount of use. When the ratio of the Grid related activity to the total network traffic grows, these differences can impede the uptake of Grid technologies. These issues can also put pressure on commercial Internet connection providers, as the Grid will weaken th e correlation between the "value" received through the network connection and the amount of traffic through it. (For more discussion of the Grid and its effects on network connectivity, see http://egee-ei.web.cern.ch/egee-ei/workshop_oct02/position%20papers/Dai%20Davies.doc).

### 2.3.4 National scientific collaborations and small companies

In the case of national level scientific collaborations, the requirements for additional investments are similar to international scientific collaborations, but the alternatives to Grid solutions are much more competitive. By establishing a small cluster and sharing the computing capacity within the collaboration, it is likely that most of the time the benefits through economies of scale are similar to the Grid. The Grid will be competitive once there exists a market for Grid service providers that are seen as reliable enough to replace dedicated clusters inside universities.

In the case of small companies, it is likely that due to the various software licensing issues the Grid itself will be difficult to market. A possible model would be an expanded ASP-model, where access to commercial software packages could be bought on demand . However, in the case of highly specialized software that represents the core competence of a small company, it is likely that until some kind of trusted computing solution is integrated into Grid, the Grid will not be widely used.

On the other end the potential interest for SMEs in being able to acquire compute services from the Grid in a very flexible and cost effective way remains a major factor of success for the Grid eInfrastructure in this domain.

### 2.3.5 Individual users

For individuals not belonging to a larger entity working with data or computing intensive problems, the motivation to use Grid services cannot be based on its ability to provide large amounts of computing capacity. An exception might be the use of specialized Grid resources, e.g. backup facilities, where the economies of scale are related to the specialization of the labour. However, services that use Grid security or knowledge of Grid related services to delegate some optimization tasks to the Grid infrastructure can be of more value. Simi larly the ability to form flexible virtual organizations on demand – and pool data sources into a single logical whole – can be of interest, especially with increasingly efficient broadband connections at home and increased use of domestic digital media, e .g. digital photography. Broadband connections together with general digital convergence will also make various health-related remote services more efficient. At the individual user level, it is likely that Moore's law – and its effect on the power of indi vidual PCs – will make the overhead of Grid solutions too large for general computational and data processing tasks. On the other hand, this same trend will make local backup -solutions proportionally less efficient, since the capacity of removable storage media is hindered by the inherently slow consensus -driven standardization process. Another aspect to consider for individual users is the difficulty of organizing individual digital archives and searching data from them – where advances in the speed of the storage devices are similarly constrained by standardization process. Thus Grid services

could provide an interesting platform for storing indexing information about the data owned by an individual user, for performance reasons.

In any case the large ma jority of today's individual web users will greatly benefit from Grid enhanced web services in particular for semantic and knowledge based information retrieval systems.

# 3 Current Practices and Achievements in Resource Access and Sharing

## 3.1 Introduction

This section summarise s the current achievements and best practises on policies concerning resource access and sharing at pan -European and International level across administrative and national domains.

The best practices and achievements include both **low-level issues in the area of resource access and sharing** including standard schemas for the core middleware services such as information indexing, scheduling and AAA – Authentication, Authorisation, Accounting mechanisms **(**including issuing, renewing and revo king certificates, physical and network security, security threat analysis, monitoring and accounting techniques) and **high-level** issues like Acceptable Usage Policies for network, computing and storage resource sharing, archiving, VO access, personal data privacy, entities registration procedures, incident handling procedures-sanctions etc.)

The resource access and sharing aspects are analysed across 3 different ax es: namely technical, administrative-procedural and legal -regulatory. The result of both the requirements, and the best practises, analyses should not provide a strict policy directive, but a complete yet flexible policy framework that could be applied in the different National Grid initiatives or efforts adjusted according to local particularities.

The latter principle is also applied in the GEANT -NREN framework, where GEANT is content with the national Acceptable Usage Policies (AUP) of each NREN and does not possess its own AUP. However, in order for an NREN to connect to the GEANT network a li st of high -level requirements must be met (http://archive.dante.net/geant/connect.html).

## 3.2 Resource Access and Sharing Schemas

### 3.2.1 Technical

In this section the technical schemas are analysed in the different areas, such as security, authentication, authorisation, accounting etc.

**Physical and network security**

Any compromise of the overall grid security will lead to a loss of confidence on the part of those who host valuable resources. Therefore all reasonable steps must be taken to ensure this does not occur. The most basic physical controls are not optional but necessary. Resources must be contained in a controlled environment (at minimum a locked room) that is safe from physical emergencies (fire, flood, electrical problems, etc) that might lead to a reduct ion in security. Archives, logs, backups, etc, must be secured. Waste disposal must not lead to a reduction in security. Those with access must be known and trusted.

Resources must be secured from unwanted traffic, typically with a firewall. Their vulne rability to intrusion must be minimized, for example by exclusion of unnecessary services, or hardening of the operating system and kernel. Updates must be applied, but themselves need to be secured. In addition, the panoply of available intrusion monitori ng, detection and prevention techniques must be used. Intrusion monitoring is well established, e.g. using SNORT, and so is intrusion detection,

e.g. using Tripwire or Aida. At present there is a trend towards intrusion prevention methods. There is, however, a lack of uniformity and coordination.

Grids must of course be protected from deliberate or recreational hacking. In this they are no different to other networked facilities. The difference is that grid computations may couple many remote resources, an d intrusion should be able to be tracked in a coordinated way. Although the tracking must have access to sensitive information, and will itself be an intrusion target, some way must eventually be found to coordinate the handling of intrusions that exploit the grid itself.

A light hand needs to be applied; *must* is a provocative word.

**Authentication**

Given the youth of grids, grid security implementations are remarkably developed, but still fall far short of maturity. There are various models for authenticati on, authorization and accounting (AAA) for the grid, but no consensus. Most grids rely on the Globus Security Infrastructure (GSI), which is based on the Public Key Infrastructure (PKI), for authentication.

PKI assumes a trusted Certificate Authority (CA) grants a grid certificate to an entity (a person, a machine or a service). The certificate is just a passport that identifies the entity to the grid. It does not authorize the entity to use resources in any way. PKI also assumes CAs issue certificate revocation lists (CRLs) that state which certificates have been revoked. Resource administrators must acquire the CRLs and bar those entities from using their resources.

The operation of certification structures by CAs involves social management as much as an ything. A CA must be known and trusted by other CAs, so personal interaction is necessary. This does not scale well globally unless the number of CAs is limited. A single CA cannot alone handle grid certification for a whole country, so devolution of autho rity to Registration Authorities (RAs) is essential. The role of the CA is to root the chain of trust, establish policies and practices, and be guardian over the root security. The CA must select RAs, who become the real certifiers of trust, so that their geographical distribution reflects that of the users.

A hierarchical root for all the national CAs has proven less workable than a forum, a Policy Management Authority (PMA), where trust can be established and maintained. The GGF is debating a global PMA infrastructure. Europe leads by example in this arena: the DataGrid CA Group has evolved to become a PMA in all but name, including members from Russia, USA, Canada and Taiwan; and the Framework 6 project EGEE will turn it into a European PMA.

The PMA will debate grid security policy. For example, within the CA Group there is a healthy and continuing debate regarding the current reliance on PKI. A crucial issue is the security of a user's private key. Whatever the mechanism, a user should not be able to compromise this by lax practices such as writing passwords in obvious places. Many CAs feel that alternative mechanisms may secure the private key more effectively than PKI, although possibly rooted to a PKI CA, and granting only short-term credentials. The propagation of trust relationships from users to resources, and policies from PMAs to resources, is a debate waiting to happen. Issues of fine detail, such as automatic mechanisms for CRL propagation, also require debate. Forward progress requires intensive debate.

**Authorization**

Current implementations of Grid security aim to keep Authentication, described above, and Authorization as two separate processes. Authentication proves a user's identity but says nothing about that user's rights to access a partic ular Grid resource. This is where Authorization comes in. The reasons for this approach include the fact that the electronic identities are issued by

independent bodies (CAs) that are not able to confirm assertions about an individual's membership of a Virtual Organization (VO) and also the fact that the possession of a single Grid identity that allows access to multiple Grids makes things easier for the user.

Early implementations of GSI in V2 of the Globus Toolkit, based on PKI as described above, contain a very primitive Authorization scheme whereby a user's Grid identity, the Distinguished Name (DN) in the X.509 certificate, is mapped to a conventional local computer account on each system in the Grid. Authorization therefore consists of the manual entry of a particular DN into the so-called grid-mapfile. Access to resources is then controlled by conventional local computer access-control based on the local username.

One of the main aims of Grid computing is to control and share resources based on membership of Virtual Organizations. The primitive grid-mapfile approach only performs Authorization for individuals and has no concept of VOs or groups of users. One early solution for this problem was the development by DataGrid of a VO management system based on a database of VO users, together with a modification to the mapfile that allows the use of a dynamic pool of local computer accounts. VO managers register the DN of bona fide members of their VO in this database, implemented as a LDAP server. Resource administrators wishing to support this particular VO then run a DataGrid tool to extract information from the VO database and automatically create entries in a grid-mapfile mapping these DNs to a pool of local accounts for that VO. This approach simplified the management of coarse-grained access control to resources on the basis of a user's membership of a VO.

The requirements for Grid Authorization include much more fine-grained access control decisions, whereby a user's possible role(s) within a VO together with his membership of particular sub-group(s) of the VO should also influence the Authorization decision. One example of an approach to address this requirement is a joint development of the DataGrid and the DataTAG projects, namely the Virtual Organization Membership Service, known as VOMS. This online service provides authenticated users with cryptographically signed authorization assertions, in the form of attribute certificates, which can then be presented to Grid services for them to make local Authorization decision.

The question of local authorization, within the Grid resource, is more complex: it has to manage the local rules and policies, in addition to the VO policies expressed in the VOMS attributes, make the authorization decision and enforce its results. An example of a current solution to this is the DataGrid Local Centre Authorization Service (LCAS), the Local Credential Mapping Service (LCMAPS) and the Grid Access Control Library (GACL).

The technical examples give above are far from exhaustive. Many projects are working in this fast moving and challenging area as can be seen by the very active working and research groups on this topic in the Global Grid Forum Security Area. Many aspects of Authorization are being tackled including requirements, frameworks, mechanisms, architectures and most importantly the discussion of standardized protocols, descriptions and interfaces for Authorization systems, particularly as related to the move to the Open Grid Services Architecture.

**Accounting**

Usage of grid resources is of interest to many parties. Consumption of resources is very important to the administrators of the resources. It is potentially a chargeable item to research funds and can server as an instrument of policy for institutions, funding agencies and governments. Up to now no cost charges have been applied in any Grid resources (which is different from the Supercomputing cases), since the accounting background is a prerequisite for billing. Accounting policies are also in

their early stages. Accounting is of key importance for the integration of super-computing centres in the Grid. No supercomputing centre or large cluster can just give resources to the Grid community without taking care of revenue for their users. As long as the demand does not saturate the resources (typical in research test-beds) the problem does not come to the surface; however when there will be lack of resources accounting and billing will be an important issue.

Accounting is also important for statistical purposes providing history parameters in grid resource markets. It may even be traded within a debt portfolio, or possibly a futures market. The accounting must eventually be treated in the same way as all other financial instruments, properly regulated according to well-understood best practices and conformance measures. Its implementation must not be as it is now, using ordinary hardware, but must avail of standard transaction processing systems typical in financial markets. This is not to say that it must be an excessive burden just that it should adhere to the norms that one would expect in financial affairs.

### 3.2.2 Administrative-procedural

The existing Grid projects have made significant progress in this area. The current practices and achievements of some of these projects are described here as input for future deliberation and work in this area. In particular a summary of current schemas and approaches of LCG and EDG projects is given before the analysis per category.

The LCG-1 security group has released six documents for Security Policies and Procedures, mainly based on the experience collected from the EDG project:

- Security and Availability Policy for LCG (Prepared jointly with GOC task force)
- Approval of LCG-1 Certificate Authorities
- Audit Requirements for LCG-1
- Rules for Use of the LCG-1 Computing Resources
- Agreement on Incident Response for LCG-1
- User Registration and VO Management

Four more still to be written (with the GOC task force)

- LCG Procedures for Resource Administrators
- LCG Guide for Network Administrators
- LCG Procedure for Site Self-Audit
- LCG Service Level Agreement Guide

The Security and Availability Policy objectives include:

- Agreed set of statements
- *Attitude* of the project towards security and availability
- *Authority* for defined actions
- *Responsibilities* on individuals and bodies

The document also insists on control of resources and protection from abuse, and tries to minimise disruption to science. Obligations to other network (inter- and intra- nets) users are also considered, It applies to Resources, Users, Administrators, Developers (systems and applications), and VOs. It does *not* override local policies. The Policy is prepared and maintained by a Security Group and GOC, approved by a forum with representation of Resource Centers, and formally owned and adopted as policy by a relevant board. The top-level policy is reviewed at least every 2 years.

Users get their individual certificate from national CA authorities and then register once for a selected Virtual Organization (Project), using the Grid certificate through a web form, and

accepting User Rules. A robust VO Registration Authority is needed to check that the user actually made the request, and is a valid member of the Project. User data is distributed to all test-bed sites.

**Billing**

As mentioned in the technical area, accounting is a prerequisite for billing and both are of key importance for the integration of super-computing centres and large production clusters in the Grid. For this reason a proper currency must be identified to express prices, so one can take into among others account prioritisation and quality of computing (e.g. 64-bit machines). The price itself can be negotiated by the market, depending on the demand vs. Supply. There could even be the option of taxes to avoid false prices on the market.

**Issuing and renewing certificates**

Authentication on the Grid requires robust procedures for establishing and confirming identity. It is often impossible and frequently undesirable to require individual users to register at each and every Grid site. The Grid identity credential, today an X.509 certificate and its associated private key, therefore plays an important role in that this is used as the primary authentication of the user. In turn, the authorization of access to resources is granted by Virtual Organizations (VOs) and granted or denied by resource owners via the association of Authorization assertions to the Grid identity credential.

The model used today in the EU DataGrid and related projects is that individual users obtain an identity credential from one of the approved Certification Authorities (CAs), which is then used to authenticate them with the Grid. The long-term aim is that this one credential can then be used as the identity basis for Authorization in multiple Grid projects and/or multiple VOs. The use of the one Grid identity across many sites and projects means that we have had to define policies and procedures of sufficient quality and robustness to be acceptable by all.

To achieve this, the DataGrid CA group has involved active participation of CA managers from several other Grid projects across many different countries. Some of the EU CrossGrid CAs were founder members of the group early in 2001, with the US DOE Grids CA joining soon afterwards. Further expansion of the group, driven mainly by the identity requirements of the global LCG project, has resulted in the approval of the remainder of the CrossGrid CAs together with Grid Canada and ASGC Taiwan, a current total of some 20 CAs. Many other national CAs have since joined the EDG CA group and are working towards approval. The current list of new CAs includes Hungary, Israel, Pakistan, Belgium, and Armenia.

The EDG CA group has defined the minimum acceptable standards for the operation of CAs and their associated Registration Authorities (RA), and for the CA/RA policies and procedures. More details are available on the EDG WP6 web (http://marianne.in2p3.fr/datagrid/ca/). When EDG ends, it is planned that EGEE will turn this CA group into a more formal European Policy Management Authority (PMA).

The experience gained building the DataGrid PKI and inter-Grid authentication with projects such as EU CrossGrid, US DOEGrids and Grid Canada was valuable input to the two GGF groups tackling Grid Certificate Policy and CA Operations issues and related work to establish worldwide trust via multiple PMAs under the auspices of the GridPMA body. It is very important that the EU CA PMA continues to work in a global context thereby allowing cross-authentication between Grids across the world.

**Revocation of certificates**

The Certificate Revocation List (CRL) is an important component of the PKI. Each CA maintains a CRL placed at a published URL. This list, digitally si gned by the CA to confirm integrity, contains the serial numbers of previously issued certificates which are now revoked and therefore no longer valid. Reasons for revocation, which are specified in the CA policy and procedure documents, include the loss o r compromise of a private key or the fact that the entity is no longer entitled to hold the certificate.

EDG and LCG have established procedures to ensure that these CRL's are updated promptly and regularly and that all sites copy them frequently.

**CA Root Repository**
An important issue in the operation of any PKI is the secure distribution of the CA certificates containing their public keys. In the Grid, these “roots of trust” are self-signed, so there can be no in-band digital signature confirming the ver acity of the information. Alternative methods of distribution have to be used. The certificates of the major commercial CAs, for example, are distributed as part of the web browser software. Today in EDG and LCG, the list of approved CAs and their public k eys are stored on project operated web servers and distributed with the Grid middleware as part of the standard software distribution mechanisms.

The TERENA task force on Authentication and Authorisation Coordination for Europe (TF - AACE) has recently created a repository for storing the certificates and policy documents of NREN CA’s. This is aimed at facilitating the use of PKI via the easy access to secure information about participating CAs. They have written a document defining the policy whereby the CA information is confirmed and securely transmitted to the repository. This is an important step forward and already several of the approved Grid CAs are included. Discussions have started for the addition of more Grid CAs.

### 3.2.3 Legal &Regulatory

**Acceptable Usage Policies (AUP)**
Keeping in mind the infancy of Grid deployment within production environments there is not enough experience on the legal and regulatory aspects of operating an eInfrastructure. At this first attempt input has been received from the relat ed experience of the research networking community (NREN AUPs and actions-sanctions in case of violations) as well as for corresponding work of the LCG Security Group on *Security and Availability Policies for LCG* . Other projects or national initiatives are welcome to participate.

In http://archive.dante.net/geant/connect.html there is a list of the AUPs of those European NRENs connected to the pan -European GEANT backbone. Although there is a var iety of AUPs covering different issues, a common basis for all AUPs can be identified. The main areas that the AUPs cover are the *eligible user communities*, their *rights and liabilities* including admissible usage and possible some exceptions or extreme cases (disasters).

**“Eligible user communities ”** for research networking infrastructures (i.e. that are allowed to connect and use the research network) are those aiming at education and/or research, rather than commercial profit and the greater "market" . Alt hough there is a variation of eligible users in the different NRENs, the user communities can be categorised in 3 rough categories:

- Education and research institutes (Academic and Research institutes & schools)
- Supporting public organisations (research min istries, libraries and some times hospitals or cultural organisations)

- Other institutes not belonging to above categories but which use the network for research or education purposes (e.g. companies or other organisations) , and usually have temporal connections.

Under "**rights and liabilities**" a long inventory of articles usually appears, stating the admissible and inadmissible (prohibited) use of the network. In case of **violation** of the above -mentioned articles and prohibited use by the connected users, the NREN has the right, with or without a proper formal warning, to perform the following actions or sanctions:

- to delete the relevant offending data (e.g. pirate copies, messages with unlawful content) with or without prior notification
- and/or to suspend a specific network service or port that causes the violation
- and/or to suspend access by the user (if feasible) or the connection of the user's organization or company, with or without prior notification to the user or his organization or company and without the user or the user's organization or company accruing any claims to damages as a result. The connection of the offending network will be restored when the latter conforms to the rules.

The LCG Security Policy specifies the following policy compliance and sanctions in case of policy violations.

**Policy Compliance**

LCG proposes that Resource Centres conduct a self -audit of their compliance with their Policy following a procedure dictated by the appropriate central GOC. Self -audits or GOC independent on-site audits will be required for the continued recognition of the service being operated.

**Legislation Compliance**

Since not all countries have uniform or consistent legislation, LCG proposes to apply policies uniformly across all sites without violating local legislation wherever possible. If this is not possible, country-specific exceptions or extensions will be made to this policy and its associated practices and procedures described explicitly in an Annex.

**Exceptions**

In exceptional circumstances LCG accepts that the emergency actions taken may violate their policies provided it is for the greater good of pursuing or preserving legitimate LCG objectives. Still the exception should be minimised, documented, time -limited and authorised at the highest level commensurate with taking the emergency action promptly, and the details notified to the GOC at the earliest opportunity.

**Sanctions**

According to the LCG policy document, resource providers and their operators or administrators who fail to comply with the es tablished policies, or its associated procedures and practices, may lose the right to have that service instance recognised by the project until compliance has been satisfactorily demonstrated again. The test of compliance will be an independent Audit. Th e same applies to the different entities such as users, administrators, developers or VOs.

## *3.3 Issues and Achievements in Resource Access and sharing*

### 3.3.1 Supercomputing centres

Major eInfrastructure resource providers, such as Supercomputing Centres, face a variety of policy-related issues. Many of the issues are shared with other resource providers (appropriate SLAs, AUPs, security approach, access capabilities, and so forth), but some issues are different because of the scale and capability of the Centres them selves. The issues that are shared with other resource

providers can be largely addressed by the same set of policies and procedures that govern resource providers, however, the issues that are different will require additional policies and considerations.

Major Supercomputing Centres have a lifelong experience in the operation of production class research infrastructures. In addition to leading state of the art computational facilities and physical resources like storage and network bandwidth, they prov ide a broad range of capabilities including research programs, development projects, advanced applications, consulting services and so forth. Supercomputing Centres are strongly focused on technology transfer from R&D in computer and computational sciences to national and international research infrastructures. Therefore, software development as well as the development and deployment of new eInfrastructure tools is often a key component of their activity, as they push technology and service levels for science and research.

There is no doubt that the leading edge Terascale computing platforms operated by the major Supercomputer Centres provide an efficient way of handling most of the tightly coupled, large scale, demanding parallel applications that are b eing deployed and run by the scientific community today. This may lead to operational models of supercomputing grids (like the DEISA project in Europe) where the emphasis is shifted from the distribution of applications across two or more platforms to the distribution of computational workload across a limited number of super -nodes, at a continental scale.

A challenging issue in a policy framework is the interfacing and cooperative operation of this kind of coarse grained distributed supercomputing infras tructure with the more traditional grids made out of a large number of lightweight computing nodes.

Besides integration and interoperation, the policy framework should also address coordination between some of the research and development activities. In some cases, various development projects will adopt different approaches, maybe even divergent views, and it will be important to balance and support these efforts until the user community determines which is best. While various types of facilities can be governed and managed by policies such as SLAs and AUPs, resources such as software, services, and maintenance of software require different types of agreements, and in many instances, these agreements will have to be negotiated on a case-by-case basis.

Some of the Centres research programs focus on pushing the leading edge of technology or in exploring new technical capabilities to a small subset of other Centres or research organizations. Other Centre research programs address leading edge science applications which only involve a restricted number of highly motivated users. In either of these cases, it will be important for the policy framework to recognize these unique capabilities and programs and provide flexibility and encourage this type of collaboration.

One approach being used by the NSF Extensible Terascale Facility (ETF) to help coordinate development activities has been to create a series of specific Working Groups with representatives from each Centre. Each Working Group develops a c harter, establishes a technical lead, and then takes responsibility to address the issues and coordination on behalf of the entire community. Another approach being used for more difficult development and leadership activities is to identify full-time positions for key coordination and integration activities to ensure that the activity receives the attention required for success. This includes activities such as network design, eInfrastructure architecture, and grid system services.

### 3.3.2 Grid projects (LCG, DataGrid, INFN)

In the last 3-4 years, several projects have been set up to test the feasibility of large computing and data oriented grids. European projects were pioneers in deploying test-beds and developing middleware to exploit these new technologies.

In 1999 **INFN** started looking at grid technology as a possible solution for High Energy Physics (HEP) computing and launched a first test-bed for the evaluation of the Globus toolkit, one of the most promising middleware packages available. This evaluation also stimulated the interest of the rest of the HEP community and led to the proposal of a European Project: **DataGrid** which was approved by the EU and started officially on the 1st of January 2001 with a 3 year contract. The **LCG** (LHC Computing Grid) followed in 2002.

The present version of DataGrid middleware is used in LCG and offers, on top of the Virtual Data Toolkit - VDT (Globus and Condor) functionalities, other services such as a Resource Broker (Workload Manager) and monitoring tools for the grid. A robust storage management and the Virtual Organisation Management System (VOMS) are foreseen to be included in LCG before the end of 2003.

The user **access policy** is regulated on the basis of the VOMS which confirms the affiliation of the user to a specific Virtual Organisation with specific rights and privileges. The user communities which have already participated in the test-bed activity belong to the High Energy Physics, Biology and Earth Observation fields. A Virtual Organisation has been defined for each community and in case of HEP experiments one VO each has allowed to separate the activity of different groups of researchers.

Local resources in many test-beds are typically insignificant, and sharing policies have not been considered as a main objective (as the test-bed word indicates). Also instability of the middleware along with scalability problems, have prevented a wider experience. However one of the obvious problems is the conciliation with local resource sharing policies in computer centres. The usual batch queuing systems (LSF, PBS, SGE, etc) in production, with a well defined sharing-/priorities-scheme for local users, have to be integrated into a grid-aware framework. Moreover, in some cases this must happen in a multi-grid environment, where several grid projects with often incompatible middleware want to access those resources at the same time. Temporary solutions have been found for compatible middleware cases (for example CrossGrid and LCG-1).

Another clear issue is accounting and its relation to quota allocation and control. While at a local level a system quota allocation check, by group or individually, is possible, in a grid framework this has to be controlled at a Collective Service Level, taking into account the actual results of accounting at local Resource Centres.

**LCG**

The LHC Computing Grid (LCG) Project was set up as a deployment project within the LHC community to take the results from the several grid middleware development projects with the following goals:

- To deploy and operate a prototype distributed computing environment for the LHC experiments,
- Learn how to maintain and operate a production grid on a global scale,
- To gain experience in close collaboration between the LCG regional centres (Resource Centres in EGEE terms),
- Focus on building a production-quality service, addressing issues such as robustness, reliability, supportability, fault tolerance and predictability rather than bare functionality,

- Understand what is needed to allow full integration of the grid service with the scientific computing services at the regional centres, moving away from the dedicated grid testbed scenario where a lot of real issues had been avoided.

To address these goals, the project scope focused on 3 main areas:

- Integrate a set of middleware from diverse sources, and coordinate and support it's deployment to the regional centres,
- Provide the operational services to enable it to be run as a production-quality service,
- Provide support services - both traditional helpdesk -style user support as well as direct assistance with integrating the application software with the grid middleware.

By the very nature of the project, many of these activities are done in the spirit of prototyping - trying various solutions and evolving them in the light of experience and feedback.

To help understand some of the **policy issues** and to broker and implement the agreements between the resource providers that it was clear would be necessary, the LCG set up a Grid Deployment Board (GDB). The members of this board are representatives of countries that have a regional centre (one representative per country), and representatives from the applications. The function of the board is to negotiate and agree modes of operation of the grid service that are acceptable to the resource providers involved, addressing such issues as user registration, security, service integration with local infrastructures and so on, while providing the applications with the functionality that they require.

The major successes of the GDB in the first year of the project have been mainly in the area of **user registration, resource access** and **security**.

**User registration**: The LCG community spans the world, with resource centres in Europe, Asia and the US. It was regarded as crucial to the success of the project that an LCG user would have to register once only to obtain access to all available resources that his Virtual Organisation was entitled to use. In addition, the funding agencies had made it very clear that there would be no discrimination between users - that every member of a VO would have access to all the resources available to that VO. The GDB negotiated an agreement that a minimal set of personal information need be collected for each user, and that that information would be collected as the user registered to use the LCG service. The US groups agreed that that information need not include sensitive data such as nationality that is normally required to use any US-DOE resources. In addition it was agreed that should the requirement for such information to be collected arise, those sites that required it would have to contact the user directly to obtain it, since it was made clear that most European agencies would not provide such information to a 3rd party. This agreement means that all LCG users are currently able to access all resources provided to their VO via a single registration. This situation could change in the future, especially as US sites become more cautious about who is accessing their resources. Currently the US sites are open scientific establishments, but should sites that restrict access wish to join the grid then work will be required to understand how that can be managed in a reasonable way.

**Security**: The LCG has a standing security task force under the aegis of the GDB. This Security Group is charged with advising the GDB on all matters related to security, including policies and procedures on User registration, Authentication, and Authorization. The group has written and GDB has approved a number of documents expounding policy for the operation of the LCG Grid. There is a top-level “Security and Availability Policy” which has 3 main objectives; it describes the *attitude* of the project towards security and availability, it gives *authority* for defined actions to certain bodies and individuals, and it places *responsibilities* on individuals and bodies participating in the project. It is important to note here that the title deliberately includes the word *availability* as

the aim of the policy is not only to prote ct resources and data from abuse, but also to promote the LHC science mission by maximising the availability and integrity of the services and data. The policy refers to a number of associated documents, these being Procedures, Rules, Guides or other technical documents required to implement the policy. To date there are 5 associated documents covering "User Registration and VO Management", "User Rules", "Approval of Certification Authorities", "Incident Response" and "Audit requirements". Documents under d evelopment include Guides for Resource Administrators, Network Administrators and Service Level Agreements and a document on Procedures for Site Self -Audit, auditing being one of the requirements of the top -level policy. The agreement of all LCG sites on t hese policy documents was seen as an important requirement for the successful operation of a large production Grid, such as LCG, and the fact that this was achieved is one of the major successes of the GDB.

**Issues of Interoperability** : LCG is already in t he situation of having to address issues of interoperation with other grids at several levels. Such issues will become more and more important as grid infrastructures develop nationally at different rates. Today, the middleware is not capable of supporting true interoperability, except at the very basic level of all grids running the same set of middleware AND using compatible operational policies. This last point is critical. It is clear that this will remain as a significant issue even as the middleware is developed to permit interoperability solutions.

In the current situation, to interoperate two grids must run the same middleware. It is unfortunate that the desire to build independent national and community infrastructures and then to peer with others far exceeds the technical level of the middleware. It will be at least a year until there are real software products that address these points. In order to arrive at such a situation there must be clear agreed definitions of grid service interfaces, agreement on protocols, and service behaviour. However, the real work will still remain at the policy level. The problems in different basic assumptions (i.e. operational policies) in recent work that LCG has done to try and interoperate with the US Grid2003 infrastructure have been seen in several different areas:

- Mapping of users to user accounts, and application assumptions and expectations
- Assumptions of IP connectivity to cluster nodes - many applications assume outgoing connectivity, but many site s cannot support this either physically or because of security policy
- Assumptions about the execution environment

The schema used to **publish information** must also be compatible from one grid to another, otherwise the discovery of resources is impossible. Within LCG (and indeed within HEP) there is an agreed standard schema for MDS - the so-called **GLUE schema** that was a joint specification by the several HEP grid projects. At a fundamental level this permits the interoperation of grid systems. However, even here with an agreed specification, it has been found that implementations of information providers and users can make differing interpretations of the specification. There is much work still to be done in fixing such schema. A more fundamental probl em arises then if two grid infrastructures wishing to cooperate use very different schema or even information systems. How resource discovery and use would occur between such domains requires much work to understand and prototype.

# 4 Requirements (link to questionnaire)

## 4.1 Introduction

This section will provide a brief list of requirements covering all the entities identified in the previous section. The list of requirements can be also used in the questionnaire that is going to be circulated on the first day of the event. The eIRG members will be requested to respond with answers from their respective national initiatives (possibly passing it to experts able to summarise their "actors" needs or after circulating the questionnaire inside their countries and analysing the answers).

After the entities' requirements' analysis and the shaping of a policy framework or architecture, it would be feasible to provide Service Level Agreements between the different entities (e.g. between end-users – providers). The requirements will be identified for each separate entity group (e.g. end-users, resource providers, etc) starting from the end-user requirements.

End-user (and others entities') requirements can be classified in the following indicative areas:

- Security including some of the relevant to Public Key Infrastructure issues:
- Privacy (or Encryption or Confidentiality), i.e. the assurance to an entity that no one can read a particular piece of data except the receiver(s) explicitly intended.
- Authentication, i.e. the assurance to one entity that another entity is who he/she/it claims to be (including Certificate Authority issues)
- Integrity, i.e. the assurance to an entity that data has not been altered (intentionally or unintentionally) between "there" and "here," or between "then" and "now."
- Non-repudiation, i.e. the assurance achieved through cryptographic methods which prevents an individual or entity from denying (in a legal sense) having performed a particular action related to data
- Reliability, robustness and high-availability
- Broadband network access translated to detailed network requirements (bandwidth & QoS reservation vs. over-provisioning, and if required other Network SLA aspects (network availability, MTBF, MTTR etc.)
- Job scheduling monitoring and accountability. The users would like to monitor their jobs (including where they have run and their statistics e.g. CPU time, storage space etc). If the resource provider will be charging them, then their expectations would be higher and accountability would be a requirement (from the provider side).
- Ubiquitous access – Mobility: Users travelling around the globe being able to submit their jobs wherever they are.

On top of the user requirements a list of further requirements will be added covering the other entities (VOs, Providers etc.). Some of the above user requirements should be valid from the other entities' viewpoint e.g. provider's viewpoint, such as security issues and broadband network interconnection of resources. More such requirements include , among others, the efficient usage of resources (planning for geographical and VO-demanded distribution, job scheduling and balancing, logging and book keeping, development of management tools). For the efficient usage of resources the experience from the power grid could be exploited (e.g. basic resource units for standard needs, burst units for special occasions, backup units for abnormal conditions etc.)

## 4.2 Entities' requirements identification

Naturally enough, grid users like to monitor the statistics of their jobs, particularly those that run for long periods. Simple things, such as where they run and their basic statistics, e.g. CPU time,

storage space, etc, but also more complex items such as bandwidth consumption, message densities, open files, swapping rate s, and any other parameters that might significantly affect the performance of the user's job. This is very suited to the web, which can support delivery to a variety of platforms, e.g. desktop PCs, mobile phones or PDAs. A good deal of this technology already exists, for example within INFN's GENIUS, but clearly it needs elaboration, both in acquisition of raw information and its analysis for presentation.

The above are important also to the resource providers and infrastructure operators, but they also need the ability to navigate through multiple levels of aggregated status down to individual users if need be. This is a multi -faceted requirement, needed for resource and network optimization, job scheduling, load balancing, for handling errant jobs, to han dle those who attempt to subvert the mechanisms to gain non-malicious advantage, and most importantly for handling security incidents. Roles and privileges must be defined, and the linkages between them established. VOs have an interest in all this, and in deed different groups of providers, operators, and their roles and privileges, may be defined as VO attributes. This generalized facility does not yet exist, although some centres have begun relevant research.

Finally this information will increasingly be a valuable input to government and EU policy, for planning for geographical and VO demand distribution and their institutional, regional, national and international, social and economic effects. As with the grid -enabling of astronomical observatories, it is likely that long -term plans will arise for special -purpose grid facilities, and perhaps even their categorization, e.g. basic resource units for standard needs, burst units for special occasions, backup units for abnormal conditions. These are not flight s of fancy, for example, an emergency grid was mooted amongst Atlantic coastal countries for the EU Interreg3 program.

If the resource provider will be charging them, then a user's expectations will be higher, and all parties, including regulatory authorit ies, will demand transparent accountability from the provider. The technology will need to support the normal double -entry balancing of financial transactions to cover, for example, resources committed but not yet delivered, resources used but not yet paid for, bad debts, asset depreciations, etc. Some of this exists in prototype form on existing grids. The IT sector will have no difficulty providing fully -fledged accounting once the economic imperative is there. The various management views of the informat ion are tantamount to database views, and the IT industry has much to teach the grid community in this area.

#### 4.2.1 User requirements

#### 4.2.2 Resource Providers requirements

#### 4.2.3 Others requirements

## *4.3 Requirements analysis*

## *4.4 Conclusion*

*Stating among others that although the above requirements have been identified from the experience of the editorial team a thorough requirements -capture and -analysis effort should be engaged in all national domains under the eInfrastructure initiative assisted by the EGEE project. In this way the questionnaire will be justified.*

# 5 Policy Framework for Resource Access and Sharing

*[To be completed at a later stage. At the first stage only the procedure to reach the desired outcome will be described]*

## 5.1 Introduction

According to the proposed methodology a policy framework for resource access and sharing at pan-European and international level and a list of policies will be drafted after capturing & analysing the multidisciplinary entities' requirements taking into account the current experience & best practices of the major Grid projects - Supercomputing centres. In addition, a roadmap for the future development of a political and administrative framework in Europe and internationally to allow a real effective exploitation of the eInfrastructures will be produced.

## 5.2 Combined Analysis of Requirements and Best Practises in Resource Access and Sharing

## 5.3 Policy Framework Proposal

## 5.4 Policy Framework Roadmap