e-IRG
e-Infrastructure Reflection Group

e-IRG White Paper 2009

Disclaimer

The e-IRG supports the production of the White Paper as a mechanism that raises new potential policy areas in e-Infrastructures. In this context the material of the White Paper is to be understood as work in progress and the views presented as reflecting in general those of the e-Infrastructure communities working on the relevant topics (which are clearly mentioned in the document) and not necessarily those of the e-IRG. The e-IRG views are reflected as such in the so called policy recommendations document.

Table of contents

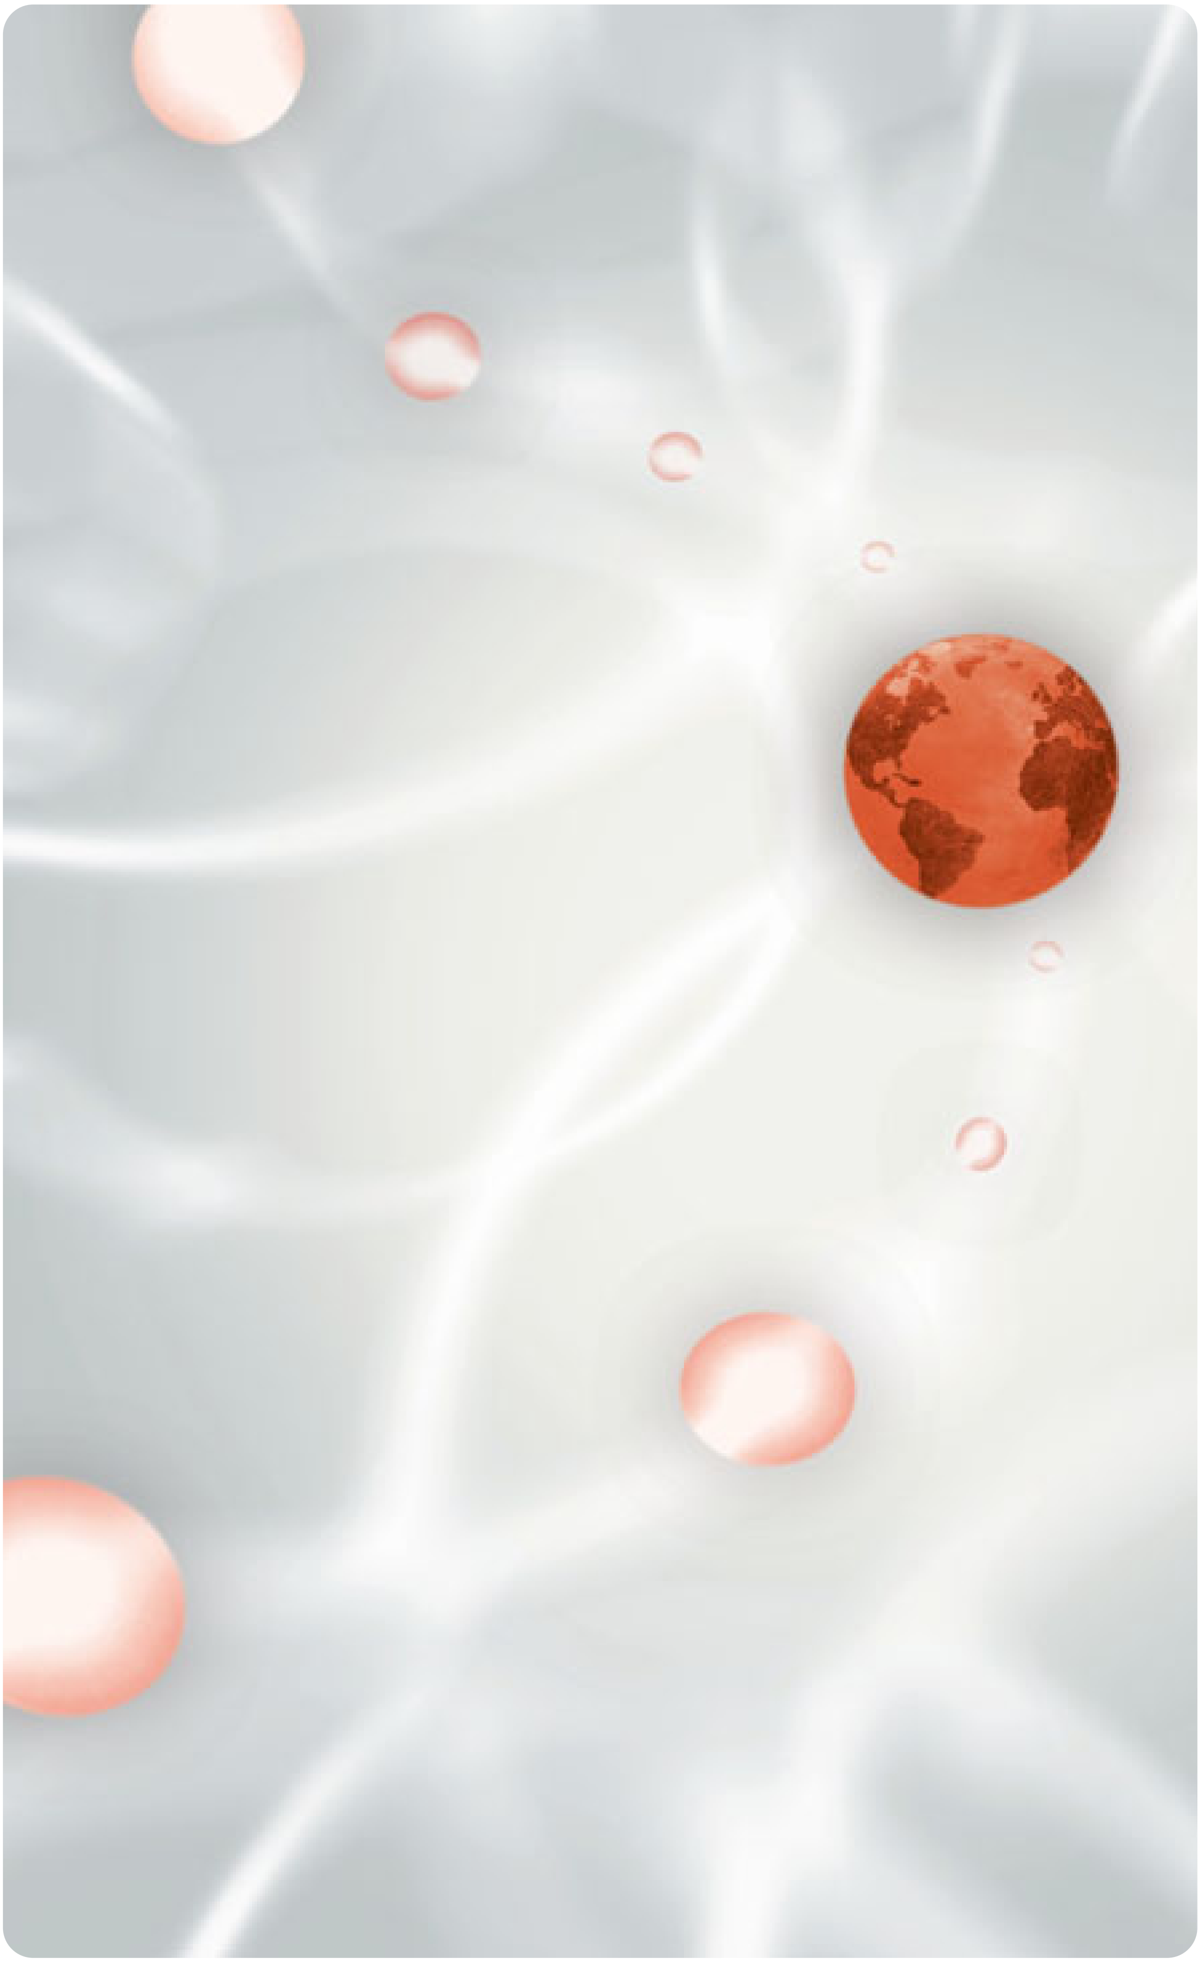

Foreword

The e-IRG White Paper is a live document summarising on-going discussions around key e-Infrastructure areas and topics that require immediate policy actions. It serves as a means to summarise analysis and dialogs by relevant stakeholders and forms the basis for proposing formal e-IRG recommendations at the national and EU levels. The topics of this e-IRG White Paper result from several rounds of reflections in a consultation process that encompassed large parts of the e-Infrastructure community. The reflections identified three key target audiences for this document:

- Policy makers on the governmental and inter-governmental levels, dealing with funding, privacy and other issues that are becoming more and more crucial due to the broader uptake of the e-Infrastructure–enabled technologies and processes.
- Service-providers on the e-Infrastructure domain, such as organisations operating research networks or computing centres – or projects that build on this base in order to provide higher-level multidisciplinary services.
- Existing and new user communities, looking for a broad overview of the capacities and capabilities that the current and near-future e-Infrastructure can provide.

The challenge of trying to serve such diverse categories of audiences – with quite different points of view and information requirements – would itself be a major one when tackling a multifaceted topic such as the e-Infrastructure. Indeed, the target is pushed even further by organisational and technological changes in and around the e-Infrastructure domain.

Nevertheless, and in order to meet this ambitious challenge, this White Paper aims to be an invitation for new high quality input from all involved organisations, communities and disciplines to join in and enrich the discussion concerning the future of the e-Infrastructure. Rather than aiming to be the last word presenting the blueprint for the future, the goal is to present, with sufficient background, the options such diverse and multidisciplinary groups must take into account in the near future about the e-Infrastructure. The future of the e-Infrastructure and its positive impact on society depend to a large extent on increasing the size, diversity and activity of the target audiences mentioned above. Translating the users' needs into coherent technological, political and organisational plans is an iterative process and this White Paper should be a powerful instrument for fulfilling this goal. We hope that this document provides tools for inclusion of larger and more diverse stakeholder groups in the decision making process.

Because of the above reasons, we are very pleased that the Swiss and French e-IRG workshops offered such a great opportunity to broaden this process. The first provided significant input to the White Paper content, while during the latter, the enthusiastic participation of the ESFRI community provided an opportunity for closer engagement of new user communities in the discussion about the e-Infrastructure. We hope that the new points of views, personal contacts and new approaches from this event form the solid foundation for reaching the key goal of the e-Infrastructure: the moment when e-Infrastructure is most often thought to mean "**everyone's Infrastructure**"!

This is aligned with the recently published EC communication on "ICT infrastructures for e-Science"[1] which "highlights the strategic role of ICT infrastructures as a crucial asset underpinning European research and innovation policies" and "calls for a reinforced and coordinated effort to foster world-class ICT infrastructures" (e-Infrastructures), "paving the way for the scientific discoveries of the 21st century".

Leif Laaksonen
e-IRG chair

1 COM(2009) 108 final: http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=COM:2009:0108:FIN:EN:PDF

Introduction

Innovation continues apace, across all aspects related to ICT and the e-Infrastructure. The goal of this White Paper is to identify strategic topics in need of policy-level action from a torrent of novel ideas, technologies and paradigms. The topics have been selected mainly by the e-IRG delegates representing their national e-Infrastructure communities, however topics have been also proposed by experts from different e-Infrastructure fields. The topics should be considered as independent, i.e. not necessarily connected to each other, mainly tackling issues which are directly or indirectly related to the e-Infrastructure as it today and its direct future. Obviously the list of topics is not exhaustive and it is not intended at covering all e-Infrastructure areas (or layers). More focus is given in areas under development where there is a greater need for policy actions; yet, in the history of previous white papers (notably this is the 6th White paper since 2003) the overall layer coverage seems quite balanced. In addition, the preparations for the next white paper (White paper 2010) will start at the end of the first semester of 2009, with topics deserving extra attention and discussion for reaching their full potential. Topics can be proposed either through the countries' e-IRG delegate, through the e-IRGSP2 support team or from consensus during the e-IRG workshops about their relevance for the progress of the e-Infrastructures.

Each topic goes through a systematic analysis according to a common structure (see also Annex II). The structure has been extensively discussed in the past during the creation of the first White papers (which were less structured). By the Austrian presidency (1st semester 2006) it was concluded that a common structure is beneficial for guiding the authors and also keeping the size manageable. In the structure agreed upon, the lead paragraph of each section provides the summary of the policy topic, focusing on the broader societal importance of the issue. This "setting of the scene" is followed by the definition of the policy area(s), goal, policy context (or policy state of the art), proposed approach (how to tackle the issue and which actions to take) and proposed recommendations. This obviously results in a certain level of detail, which is felt to be necessary for conveying the main ideas not just to policy makers but also to e-Infrastructure service providers and user communities. A more policy-oriented document, focusing on policy-makers and e-IRG delegates, is planned as a follow-up of the White paper. This is the so-called White paper Recommendations document[2], which covers one page per topic with distilled and worked-out recommendations.

The selection of topics for this White Paper was based on the recently re-articulated mission of the e-IRG being "to pave the way towards a general-purpose European e-Infrastructure". *Global Collaboration* focuses mainly on the organisational aspects of the European e-Infrastructure, offering a global view of the bridging effect of European collaborations with the rest of the world and their strategic importance for mutual benefits, for Europe and all other regions. *Education and Training in the Use of e-Infrastructure* is a summary of the findings and recommendations of the report of the e-IRG Education and Training Task Force. *Cloud Computing* and *Virtualisation* illustrates the constant need for awareness of the new opportunities provided by innovative technologies that can be integrated into the overall e-Infrastructure. The fact that these technologies are mostly promoted by industrial entities offering a portfolio of commercial services crossing some paths of established e-Infrastructures gives a certain level of urgency for coordination of actions among related stakeholders. *Security, a holistic approach* examines security from different angles ("as a whole") covering all e-Infrastructure aspects in an effort to minimize overlaps and promote synergies among different e-Infrastructure service providers (e.g. network and grid providers) for minimising a

2 An example of such document can be found in http://www.e-irg.eu/images/stories/publ/finnishpresidency-recommendationsanddecisions.pdf

problem concerning all e-Infrastructures and all users. *Remote instrumentation*, although a "new kid on the block" is an appealing topic highlighting the role of the e-Infrastructure in handling remotely Research Infrastructures such as telescopes, sensors and remotely operated vehicles. Lastly, *Sustainability of computing-related e-Infrastructures* is focused mainly on the organisational aspects of European computing components of the e-Infrastructure aiming at providing a sustainable and seamless service to researchers in Europe and beyond.

The White Paper preparation process was steered through a series of documents, the most important of which being the Terms of Reference (ToR), the White paper Work Plan, the White Paper topics and contributors spreadsheet, the White Paper skeleton which evolved into the actual White paper, and the White Paper workshop- and contribution-templates. The WP ToR, Work Plan and structure are included as Annexes to this document. The White Paper ToR, workplan and templates (introducing the current structure) were presented and approved during the Zurich e-IRG workshop on 25 April 2008, while the White paper topics and skeleton on 27 June 2008 in Lugano. Then, on the next two e-IRG meetings in France (Paris, 22 October and 8 December), the White paper versions were presented and received a series of guiding comments from the e-IRG delegates.

Fotis Karagiannis
White Paper editor/ e-IRGSP2 deputy director

Terminology

"Capability computing" refers to serving at one single moment in time a coarse number of specialised computing tasks requiring an extremely powerful and tightly integrated computing system;

"Capacity computing" refers to serving an extremely large number of parallel tasks on a large scale computing infrastructure;

Cloud computing (or simply "Cloud") is an on-demand service offering a large pool of easily usable and accessible virtualized resources (such as hardware, development platforms and/or services) in a pay-per-use model. Clouds are usually commercial and currently use proprietary interfaces;

e-IRG Education and Training Task Force (ETTF) – a task force created in the framework of e-IRG investigating and proposing strategic actions in the area of education and training in e-Infrastructure and its curricula. A related report has been published available at the e-IRG site;

e-IRG Sustainable e-Infrastructure task force (SeITF) - a task force created in the framework of e-IRG investigating and proposing strategic actions in the area of e-Infrastructure sustainability. A related report has been published available at the e-IRG site[3];

e-Science – the invention and application of ICT-enabled methods to achieve new, better, faster or more efficient research, innovation, decision support or diagnosis in any discipline. It draws on advances in computing science, computation and digital communications;

electronic Very Long Baseline Interferometry (eVLBI) is the process of using high speed networks to connect radio telescopes separated by large distances (100-1000s of km) instead of the traditional method of recording onto magnetic tape and shipping the recorded data to a central correlator. It is part of a collaboration of the major radio astronomical institutes in Europe, Asia and South Africa performing high angular resolution observations of cosmic radio sources;

Education and Training Community Group (ET-CG) – is an OGF group aiming at bringing together practitioners in grid-related education and training (E&T) to share and develop best practice, to stimulate greater investment in Grid-related E&T and above all to build a mutually supportive community of grid trainers and educators;

European Grid Initiative (EGI) is the next phase of the implementation of capacity computing in Europe. EGI will manage no compute power, but guarantees transnational access to data and services;

Federated e-infrastructure Dedicated to European Researchers Innovating in Computing network Architectures (FEDERICA) - is European project co-funded by the European Commission aiming at implementing an experimental network infrastructure for trialling new networking technologies. This infrastructure is intended to be agnostic to the type of protocols, services and applications that may be trialled, whilst allowing disruptive experiments to be undertaken. The aim is to develop mechanisms that will allow such experiments to be run over existing production networks without adverse effect[4];

Grid is a system that federates, shares and coordinates distributed resources from different organisations which are not subject to centralized control, using open, general-purpose and in some cases standard protocols and interfaces to deliver non-trivial qualities of service. The Grid is used by Virtual Organisations, i.e. thematic groups of users crossing administrative and geographical boundaries;

International Collaboration to Extend and Advance Grid Education (ICEAGE) is a Specific Support Action project co-funded by the European Commission aimed at stimulating and supporting academic teaching in distributed computing;

3 http://www.e-irg.eu/index.php?option=com_content&task=view&id=38&Itemid=37

4 www.fp7-federica.eu

IPR - Intellectual Property Rights – refer to the controlled right of use of created items, so that the creator benefits from that use. Intellectual Property is broken down into several types, which apply to different created items: copyright, designs, patents, trademarks, protection from passing off and protection of confidential information;

International Thermonuclear Experimental Reactor (ITER) – is a joint international research and development project that aims to demonstrate the scientific and technical feasibility of fusion power. Fusion is the energy source of the sun and the stars. On earth, fusion research is aimed at demonstrating that this energy source can be used to produce electricity in a safe and environmentally benign way, with abundant fuel resources, to meet the needs of a growing world population;

Open Grid Forum (OGF) – is an open community committed to driving the rapid evolution and adoption of applied distributed computing;

Partnership for Advance Computing in Europe (PRACE) – is a preparatory initiative aiming at implementing 3-5 petaflop supercomputing systems in Europe. PRACE manages extreme computing power and a selected set of highly specialized services;

Research and Education Federations (REFEDs) wish to address the need of existing and emerging e-identity federations operating in the field of education and research in Europe, America and Asia to collaborate on policy issues. Collaboration in this area is needed to learn from each other, to align policy in order to facilitate inter-federation work, and to reach outside the research and education community to enable federated access in the larger world[5];

Repository - a storage place for digital resources. Users can easily search, access and use resources collected in a repository via an online network. A digital library is a type of repository (a particular application);

TERENA Task Force on Computer Security Incident Response Teams (TF-CSIRT) is a TERENA task force that promotes collaboration between CSIRTs at the European level, and liaises with similar groups in other regions[6];

t-Infrastructure (or training infrastructure) is an e-Infrastructure adapted to the needs of education, trainers and students. Shared t-Infrastructure would be usable by students and teachers throughout the ERA, providing easy access to educational exercises running on (good emulations of) e-Infrastructure;

Training - a short-term process to develop specific skills in a certain technical area;

Virtualisation refers to the decoupling of the resource used to provide services from the hardware the service is running on, by providing an additional abstraction layer that hides the underlying technology and devices;

White paper - an authoritative report that often addresses problems and how to solve them. White papers are used to educate readers and help people make decisions. They are often used in politics and business[7]. Commission White Papers are documents containing proposals for Community action in a specific area. In some cases they follow a Green Paper published to launch a consultation process at European level. When a White Paper is favourably received by the Council, it can lead to an action programme for the Union in the area concerned[8];

Worldwide LCG Collaboration (WLCG) - is a global collaboration of more than 140 computing centres in 33 countries. The mission of the WLCG project is to build and maintain a data storage and analysis infrastructure for the entire high energy physics community that will use the Large Hadron Collider (LHC) facility at CERN.

5 http://www.terena.org/activities/refeds/

6 http://www.terena.org/activities/tf-csirt/

7 Source Wikipedia: http://en.wikipedia.org/wiki/White_paper

8 Source Europa Glossary: http://europa.eu/scadplus/glossary/white_paper_en.htm

Executive summary of the e-IRG White Paper

Through its White Paper the e-Infrastructure Reflection Group (e-IRG) invites policy makers, e-Infrastructure service providers and user communities to participate in the discussion over the immediate future of e-Infrastructures. The document presents key e-Infrastructure issues and topics requiring policy action at the national and EU levels. The purpose of the e-IRG White Papers is to provide a "snapshot" of organisational, technological, economical, and social developments in the e-Infrastructure domain. The White Paper outlines the opportunities and challenges stemming from these developments, and proposes policy-related recommendations for the next steps, aiming at maximising benefits from the considerable investments in the e-Infrastructures in support of research and education in Europe and beyond.

Seven topics are examined in the e-IRG White Paper 2009: Global Collaboration, Education and Training in the Use of e-Infrastructure, Grid and Cloud Computing, Security - A Holistic Approach, Service-centric e-Infrastructures through Virtualisation, Remote Instrumentation, and Sustainability of the Computing-Related e-Infrastructure.

Global collaboration and Education – Training are generic areas where policies promoting a coordinated action of multiple-stakeholders along with better organisation are required. In addition, these areas clearly contain a social aspect, and mostly a technological background. In Grid and Cloud computing, technological developments play a key role in the assessment of policies, with economical issues remaining at a prominent position. Security is in general driven by technological issues too; however the main effort in this White Paper concentrates on organisational policies, advancing collaboration of the several related e-Infrastructure service providers. Moving to more specific issues, Virtualisation and Remote Instrumentation are currently dominated by technological developments; yet economical issues in Virtualisation and organisational and social issues in Remote Instrumentation will be getting more attention in the near future. Sustainability of the Computing-Related e-Infrastructure deals mainly with organisational issues of the capability and capacity e-infrastructure components, requiring policy actions for their smooth interoperation and interaction for the benefit of their users.

Entering deeper into the heart of the topics, the need to expand Global Collaboration activities finds its *raison d'être* into the broadening of the user communities from scientific domains with well-established institutional structures to disciplines with more fluid collaboration models. These new global collaboration activities pose additional challenges on the policy level. The absence of a clearly defined, hierarchical organisation to collect and prioritise the needs of the user communities means that it is necessary to establish other venues for exchanging information about best practices and future plans on the global scale. The e-IRG intends to play a major role in this information exchange.

Education and Training in the Use of e-Infrastructure is another important issue addressed by the e-IRG White Paper 2009. Taking into account the continuously expanding number of research activities supported by e-Infrastructures, learning how to best use them should be an integral part of most curricula in nearly all disciplines. The inclusion of this intellectual material on top of the disciplines currently making use of e-Infrastructures will dramatically increase the benefits resulting from the uptake of the e-Infrastructure services in new domains. The White Paper builds on the work of the e-IRG Education and Training Task Force (EEFT) to propose a plan of actions. Among other recommendations, it calls for an ERA-wide harmonisation and standardisation of distributed computing-related knowledge and skills, in accord with the Bologna process[9].

9 http://ec.europa.eu/education/higher-education/doc1290_en.htm

Cloud computing and Virtualisation, are among the most promising and innovative ICT technologies. Developments in these areas need to be taken into account when re-assessing the future of e-Infrastructure and related policies. In the field of distributed computing, cloud computing represents a new paradigm. Seeking a way to combine best aspects of the Grid and Cloud Computing requires resolving technological and policy level conflicts between two paradigms that have started based on very different requirements, use cases and basic assumptions. The Grid approach is based on requirements that were both quantitatively and qualitatively very challenging, necessitating a solution for managing remote resources that are not controlled by any single organisation in a secure and comprehensive manner. However, the complex, resource intensive computational tasks and workflows needed by the users has led to relatively complex technical and organisational interfaces. These interfaces are not optimally suited for simple use cases that are essentially aiming at replacing a single server or a local cluster with remote ones. In the commercial domain addressing this more straightforward scenario has led to the development of the Cloud paradigm, which provides simple, commodity services for on-demand provision of IT infrastructure. The analysis of the state of the art of both of these computing paradigms indicates that the Cloud offerings are not yet sophisticated enough to support complex “Grid-like” use cases directly; therefore, it is important to study both opportunities provided by the integration of cloud-based offerings into research e-Infrastructure and by promoting the reuse of the results of the Grid initiatives in the Cloud context. In other words, the e-IRG believes that the long-term interests of the research community would be best served by increasing access to and use of a mixture of grid- and cloud-based services and technologies, and encourages the integration of cloud-based services into the existing e-Infrastructure.

The Security approaches of the different e-Infrastructure service providers contain fundamental differences in the basic assumptions relating to what is considered an acceptable level of security – both from the technological and organisational points of view. Furthermore, not only are the assumptions different, but also the language used to describe and communicate them varies depending on which e-Infrastructure component is under study. Thus, in order to achieve and retain an acceptable level of overall security in a cost-effective manner, the e-Infrastructure community needs to develop common ways to specify and measure security across network, Grid, supercomputing and data domains. This will allow for establishment of cross-organisational structures for security incident management – mirroring the CSIRT collaboration between the NRENs – to provide both reactive and proactive security solutions and services for the whole span of e-Infrastructure layers.

In addition to emerging data-related initiatives, Virtualisation and Remote Instrumentation add new technological and organisational aspects to e-Infrastructure services. Virtualisation covers its first steps, especially in the research world. However the growing interest on the technology, mainly due to the economical benefits that brings along, advocate for some first policy actions: integration of virtualisation technologies in the current research e-Infrastructure and encouragement of open standards are such initial actions. The current multitude of interfaces from different vendors may otherwise endanger the e-IRG vision of an open e-Infrastructure that allows optimal use of all electronically available resources.

The e-Infrastructure–related opportunity in the Remote Instrumentation area is the ability to leverage proven collaboration models used to share “pure” ICT resources into the laboratory environment. By successfully anticipating emerging collaboration needs in this domain, the service providers can offer a proven set of tools for this new user community. The development and spread of remote instrumentation techniques and technologies that allow, remote and shared access, for its part, opens up new opportunities for scientific communities. Further research on these technologies should be supported, as should their progressive integration into the e-Infrastructure framework. Fair access, integration with Grid middleware and work on standards are the key policy issues in this area.

An issue crosscutting all of the above areas is the sustainability of e-Infrastructure. If relevant communities do not trust that services will be available in the future, the issues of integrating new components into the existing e-Infrastructure and improving security cannot be approached in a stable manner; nor is it possible to harvest the full benefits from investments in education, training and global collaboration efforts. The goal of self-sustainable funding of European e-Infrastructure services requires careful analysis of organisational and funding models embraced by all stakeholders – both resource users and providers. The models and experience of existing NRENs can act as a guideline for the computing components of the e-Infrastructure, as well as for the future data-related initiatives. Linking the proposed new capacity and capability computing organisational paradigms into with the existing NREN model and future data-related initiatives is an important challenge for the e-Infrastructure community.

1. Global collaboration

Editors: Marie-Christine Sawley <sawley@cern.ch>, Ana Bela Dias <dias@nwo.nl>,
Contributors: Bill Feiereisen <bill@feiereisen.net>, Fotis Karagiannis <fkara@aueb.gr>, Kees Neggers <Kees.Neggers@surfnet.nl>

It has been long recognised that e-Infrastructures are not restricted to country borders; on the contrary they must be planned and set up as global infrastructures in order to create an effective and competitive scientific ecosystem. The engagement in worldwide collaborations between the various e-Infrastructures will increase cross fertilization of novel ideas across large scientific communities and harmonise policies and best practices between trans-continental large scale e-Infrastructures towards the worldwide Knowledge Society. In the short term this can be achieved by supporting cooperation at global level of existing communities, such as bio-medical sciences, environmental sciences and high energy physics, while preparing for the extension of these working models to other areas. In the long term, collaboration between trans-national e-Infrastructures will not only be an extra incentive for innovation and harmonisation of the e-Infrastructures themselves but should also create superior economies of scale by facilitating seamless access to resources worldwide and by accelerating the deployment of high quality services beneficial to research and education. Broader socio-economic benefits will also emerge, through opportunities for more efficient technology transfer and for harnessing larger innovation potentials of the global networks supported by the e-Infrastructure.

1.1. Definition of policy areas

Harmonisation of major policies across continents is a key factor for ensuring superior cost-effectiveness and faster development of all e-Infrastructures. This can be achieved through an open dialogue between all parties involved worldwide in setting up and coordinating e-Infrastructures by sharing and discussing best practices and development trends. Such best practices pertain to e-Infrastructures (resources, software tools and services) supporting a wide range of domains such as large-scale data repositories, networking, supercomputing, grid computing, interactive large scale visualization and complex workflows for data analysis. Services should aim to hide the complexity of the resources while providing ubiquitous access and enhanced usability; software tools such as programming languages and software suites should exploit the full potential of the resources.

1.2. Goal of the policy analysis

Collaboration for innovation and harmonisation of e-Infrastructures worldwide is the main goal of this topic with multiple advantages regarding the e-Infrastructures themselves, being the ease of use and the resulting socio-economic impact. Open collaboration with countries outside Europe will support European policy makers in their decisions for new e-Infrastructure developments and applications. Scientific collaborations between European and non-European researchers will be facilitated, promoting cross-fertilisation of new ideas and allowing for easy access to e-Infrastructures outside the European ecosystem. Though in some scientific fields this is not new and has almost become common practice, the same situation should be made available to all researchers, independent of their research area. Global collaboration will also add value to the European e-Infrastructure by increasing the efficiency and competitiveness of the European scientific community and by favouring innovation and the resulting socio-economic impact.

1.3. Context

One of the first challenges is to decide which parties outside Europe will bring the largest contribution to the foreseen goals. The USA and Japan are at the top of the list of countries that have a mature e-Infrastructure. Other countries that can offer new perspectives due to their emerging technologies and fast growth are China and India, the last one especially as an HPC technology developer. Identification of the most beneficial organisations in Japan, Australasia, China and India for setting up collaborations can be more challenging than in the case of the USA due to possible hindered communication. Latin America, South Eastern Europe and Mediterranean countries should also not be left out due to the potential presented by countries within those regions. Also some large European projects, fundamental for the development of science and industry in Europe, which strongly depend on e-Infrastructures for achieving their full potential, should be included in the list of potential collaborations.

In this section some important organisations and their mission statements are recorded. This list could be much longer, but should be considered as an initial attempt to identify organisations with large policy impact in their regions.

In the USA there are multiple influencing organisations. The Computing Community Consortium – CCC (NSF) – was created by the Computing Research Association (CRA) in 2006 to catalyse the computing research community to debate more audacious research challenges; to build consensus around major, long-term research directions; to articulate research agendas; to evolve the most promising visions toward clearly defined initiatives; and to work with funding organizations to move those initiatives toward funded programs. CCC is funded under a cooperative agreement from the National Science Foundation (NSF) and aims at creating initiatives such as the NSF Information Technology Research (ITR) program and the NSF Cyber-enabled Discovery and Innovation (CDI).

One interesting area that CCC is focusing on is the promotion of revolutionary, ground-breaking, bold proposals that support current computer science research, as it appears that the most successful proposals are now incremental and evolutionary, lending concern that the pipeline of new ideas leading to the next big technology changes is not full.

The Office of Cyberinfrastructure – OCI (NSF) – coordinates and supports the acquisition, development and provision of state-of-the-art cyber-infrastructure resources, tools and services essential to the conduct of 21st century science and engineering research and education. OCI supports cyber-infrastructure resources, tools and related services such as supercomputers, high-capacity mass-storage systems, system software suites and programming environments, scalable interactive visualization tools, productivity software libraries and tools, large-scale data repositories and digitized scientific data management systems, networks of various reach and granularity and an array of software tools and services that hide the complexities and heterogeneity of contemporary cyberinfrastructure while seeking to provide ubiquitous access and enhanced usability.

The Advanced Scientific Computing Research – ASCR (DoE) – aims at discovering, developing, and deploying the computational and networking tools that enable researchers in the scientific disciplines to analyze, model, simulate, and predict complex phenomena important to the Department of Energy. DoE and ASCR support the Scientific Discovery through Advanced Computing (SciDAC), Innovative and Novel Computational Impact on Theory and Experiment (INCITE) and the Multi-scale Mathematics Initiative programs.

Open Grid Forum – OGF - is an open community committed to driving the rapid evolution and adoption of applied distributed computing. Applied distributed computing is critical to developing new, innovative and scalable applications and infrastructures that are essential to productivity in the enterprise and within the science community. OGF accomplishes its work through open forums that build the community, explore trends, share best practices and consolidate these best practices into standards.

In Japan, the Cyber Science Infrastructure – CSI (NII) – is the foundation for Japan's scientific research and education activities aiming at promoting the strengthening of international competitiveness in order to maintain cutting-edge scientific development. CSI acts as a responsible entity for the NREN, Grid initiative, data repositories, middleware research and development of new HPC architectures, under the National Institute of Informatics.

In Australia, the National Computing Infrastructure is a joint venture from the Commonwealth Department of Innovation, Science, Industry and Research, and the Australian National University in Canberra. It operates the national high-end computing facility, coordinates the access to specialised computing resources for specific disciplines and runs an access programme – including software development and training.

In China, the Chinese Academy of Sciences – CAS) – is responsible for undertaking government-assigned projects with regard to key S&T problems in the process of social and economic development; initiating personnel training; and for promoting China's high-tech enterprises by active involvement in these areas. High-technology projects are addressed by the central government as 5-year plans. The last 5-year plan runs from 2006 to 2010. The Chinese Academy of Sciences is the highest national advisory board for S&T.

In the greater Asian region the EC is promoting e-Infrastructure bonds with Europe through funding projects in the area of networking (such as the TEIN series, currently TEIN3 which connects also Australia and Japan) and Grid computing (such as EuChinaGrid and EuAsiaGrid).

In India the Centre for Development of Advanced Computing – C-DAC – aims at continuing to create and deploy the finest talent for further expanding the frontiers of High Performance Computing and Communication Technologies and its applications. C-DAC is under the Department of Information Technology (DIT) from the Ministry of Communications of Information Technology. Similar to the above e-Infrastructure bonds between EU and India are supported by the EC exploiting the Indo-European programme to link the local NREN to GEANT, as well as funding Grid projects like EuIndiaGrid.

In Latin America, CLARA – Cooperación Latino Americana de Redes Avanzadas – is constituted by 17 Latin American countries and its Assembly – where each country has a representative – meets every six months to define courses of action and the policies to be implemented. In addition, the EELA project series are also an important link to the distributed computing e-Infrastructures in Latin America.

The countries around the Mediterranean Sea, EU and non-EU, have been collaborating for many years in the area of e-Infrastructures. The Eumedconnect project created a research network linking the north African and Middle East countries (in particular Algeria, Egypt, Jordan, Morocco, Palestinian Authority, Syria and Tunisia) with the rest of Europe (besides Cyprus and Malta who joined the EU). Following this first step which made the first towards policies harmonisation, the EumedGrid project extended the efforts towards distributed computing producing related white papers and MoUs. The Eumed events[10] (in 2007 and 2008) act as a main enabler of collaboration with high-level representations, and policy aspects are in the centre of attention.

The recent international cooperation EC call focusing among other regions on China, India and the Mediterranean shows the strategic important of these regions.

In South Eastern Europe, a series of regional e-Infrastructure projects have been developed and executed with great success, initially in the networking area (SEEREN series) and then in the Grid area (SEE-GRID series). The main objective of such collaborations is to bridge the digital divide and integrate the non-EU countries to the European initiatives, along with their policies.

1.4. Proposed approach

Global collaborations are becoming an important and urgent issue for worldwide harmonisation and development of e-Infrastructures. The following prompt actions are proposed:

- Initial interactions between the European and non-European organisations involved in e-Infrastructure policy should be achieved through mutual briefings (e.g. CCC presentation at the e-IRG workshop in Zurich on April 2008); invitations to key contact persons from the above organisations or projects should be encouraged; the e -IRG workshops can accommodate talks and fruitful discussions on mutual cooperation and cross-fertilisation;
- Initial interaction should be followed up by regular and formal contacts. Setting up a specific e-IRG Task Force should be investigated. The Task Force would be in charge of not just investigating new possible collaborations but also actively engaging in the collaborations and regularly reporting on the results achieved to the e-IRG forum during its meetings;
- Making policy-makers aware of global collaboration recommendations;
- Involving EU members in collaborations to support e-Infrastructure directives in the context of FP programs.

1.5. Recommendations[11]

- The e-IRG acknowledges the global nature of Science, Research and Education and recognises the strategic importance of Global Collaboration between worldwide parties involved in e-Infrastructures as a key element for the global competitiveness of the European e-Infrastructures;
- The e-IRG recommends that Global Collaboration should move from its ad-hoc character to a more structured and continued mode adequately supported by the EC and national funding agencies; Europe should continue to take a leading role in global collaborations.

10 http://www.terena.org/activities/development-support/eumedevent2/

11 All recommendations are part of a separate document called “Recommendations 2009” document

2. Education and Training in the use of e-Infrastructure

Editor: Malcolm Atkinson <mpa@nesc.ac.uk>,
Contributors: David Fergusson <dfmac@nesc.ac.uk>, Elizabeth Vander Meer <evmeer@nesc.ac.uk>, Jo Newman <jo@nesc.ac.uk>,
Reviewers: Neil Geddes <n.i.geddes@rl.ac.uk>, Rosa Badia <rosa.m.badia@bsc.es>

ICT and e-Infrastructure already prove indispensable in academic and commercial spheres. The return on investment in ICT and e-Infrastructure in both contexts depends on educating European citizens so that they can exploit e-Infrastructure's potential throughout their careers and in their communities. The rapid development of readily available computational services, rich data resources and digital data collecting devices means that there are an immense number of new opportunities for innovation opening up every year. Adroit exploitation of these opportunities, in business, medicine, engineering, government and virtually all walks of life depend on the ability to understand the potential and recognise the opportunities. That adroit exploitation is essential for competitive Europe and for wellbeing in an ageing society.[12] It can only be achieved if citizens have been prepared during their education. Therefore we recommend that the EU and its Member States invest in developing educational capacity and in adapting its existing curricula to enable a wider range of graduates to exploit e-Infrastructure and to appreciate its value to society.

To deliver success, once the opportunities are recognised, requires a cohort of engineers, informaticians and other specialists. Therefore, commensurate attention should be given to educating and training these technical specialists.

Relevant education and training programmes exist in many European countries.[13] However, investment is needed to make these much more widespread across disciplines, across institutions and across all geographic regions. This can be achieved more economically through cooperation among Member States. Such cooperation should result in mutually recognised curricu-

12 http://ec.europa.eu/employment_social/knowledge_society/index_en.htm

13 http://www.iceage-eu.org/v2/msc%20courses.cfm, http://www.iceage-eu.org/v2/other%20summer%20schools.cfm

la and qualifications which facilitate worker mobility and multi-national teamwork. Policy measures are needed to accelerate this collaboration and to make it more effective.

2.1. Definition of policy areas

The policy areas discussed here relate to coordination on curricula development and the sharing of resources, knowledge and experience in order to promote e-Infrastructure education and training across Europe in the context of existing educational programmes. Policy decisions in these areas affect education and training policy makers, funders, tertiary education institutions, teachers and students. The effects of the new policies should percolate into secondary education as a consequence of their impact in tertiary education.

2.2. Goal of the policy analysis

The primary aim of this policy analysis is to induce changes in tertiary education in order to generate a sufficient stream of graduates with the relevant knowledge and skills, who are already alerted to the potential of the new technologies and the ways in which they will transform their discipline. Many studies have identified overall skills and knowledge shortages in the use of computing systems. The EC ICT Skills Monitoring Group compiled specific data on skills shortages in a report that provides details of the crisis.[14]

The increased production of graduates adept at exploiting the advances in ICT in their discipline is a primary goal crucial to the competitive success of Europe. This depends on achieving two sub-goals:

1) The injection of relevant intellectual material into the majority of curricula across nearly all disciplines and;
2) The rapid adoption of these new curricula across the great majority of tertiary education institutions in all geographic regions of the EU.

A sufficient stream of experts capable of designing, engineering, delivering and managing e-Infrastructure in all its application contexts is the secondary goal. A broad view of the forms of e-Infrastructure and the contexts in which it can be deployed must be taken when developing the educational provision for these experts. It should, in a manner parallel to that described above, influence the informatics, engineering, computational sciences and mathematical sciences departments of tertiary education institutions in all geographic regions of the EU. The curricula here will evolve rapidly as these disciplines advance to consider global-scale ICT. From this cohort, researchers will arise to develop the foundations of future generations of e-Infrastructure.

2.3. Context

A number of leading research centres and tertiary education institutions run Doctoral programmes and Masters courses that teach specific aspects of computational science, e-Science, e-Research or mathematical sciences and address the relevant educational goals. They are a significant source of the seeds from which the new curricula and educational programmes will grow. However, they are by no means as widespread, either geographically or across disciplines, as is necessary. Even where a discipline has an active programme, it is rarely engaging the majority of students.

14 http://ec.europa.eu/enterprise/ict/policy/doc/e-skills-forum-2004-09-fsr.pdf

The EU ICT projects concerned with e-Infrastructure run training programmes and summer schools[15], as do some national and regional organisations. Again, these are a vital source of experience from which to develop education and training, but they do not approach the volume of influence that is required. The following sections provide a more detailed sketch of the education and training landscape, considering curricula development and sharing of resources.

2.4. Proposed approach

Curricula Development

The majority of existing work on curricula has been achieved by independent efforts for specific courses. The ETTF (e-IRG) and the ET-CG (OGF) supported by EGEE and ICEAGE organised discussions on curricula[16] leading to the conclusions that:

1) A pervasive introduction of thinking skills into undergraduate courses that developed judgement and problem-solving in the context of modern ICT and e-Infrastructure would be extremely valuable to graduates and their employers;
2) Courses should be available to students in the final years of first degree programmes and in Masters programmes so that they can learn how to use, and in some cases develop, advanced applications pertinent to their discipline; and
3) Advanced courses should be available, bringing together informatics and application skills for those who wish to become experts in developing and exploiting the new technologies.

Access to e-Infrastructure for training

The provision of training infrastructure facilitates education and training in the use of e-Infrastructure. European e-Infrastructure must provide the facilities for essential hands-on experience. In practice, this Infrastructure is the computing equipment, digital communications, software, data and support staff needed to teach a course. Development of an EU-wide approach would advance the sharing of curricula, qualifications and teaching methods. Policy in this area should eventually resolve issues such as:

1) Access rights and their management for teachers and learners;
2) Fair sharing of costs across institutions, educational funding bodies and Member States;
3) Procedures inducing good behaviour when using shared e-Infrastructure for teaching, for example, policing misuse and encouraging submission of improvements back to the shared repositories.

One aspect of good behaviour is the recognition of contributions and the respect of intellectual property.

Resource sharing should include the use of and access to repositories of educational and training materials. Intellectual Property Rights (IPR) issues arise in this context. Resolving such IPR issues is required in order to spur greater uptake of e-Science. Currently, no specific IPR model has been widely accepted across the European Research Area (ERA).The 2001 EU Copyright Directive is an attempt at harmonising copyright law among Member States, keeping in mind certain modern requirements of the information society.[17] But the challenges arising from the sharing involved in use of e-Infrastructures are still in the process of being unravelled and addressed. There are repository examples that can be held up as models, showing the way forward, and these include the ICEAGE Digital Library and EGEE Repository, both of which provide educational materials and tackle the issue of IPR using deposit agreements and Creative

15 http://cordis.europa.eu/ist/, http://cordis.europa.eu/fp7/ict/home_en.html, http://www.iceage-eu.org/issgc08/index.cfm

16 http://www.alphagalileo.org/nontextfiles/ettf_long_report_final_july08.pdf

17 http://eur-lex.europa.eu/smartapi/cgi/sga_doc?smartapi!celexapi!prod!CELEXnumdoc&lg=EN&numdoc=32001L0029&model=guichett

Commons licences.[18] The CoreGrid repository is another example supported by a Quality Assurance process. Concerted action by the Member States can be productive in resolving access to educational licences for commercial software, data and systems.

2.5. Recommendations

- The e-IRG recommends that the national, regional or European level of investment in e-Infrastructure education should be balanced to the investment that is going into e-Infrastructure provision. This may be achieved by embedding e-Infrastructure education at undergraduate level, as well as developing curricula at postgraduate level to improve exploitation of e-Infrastructures;
- The e-IRG recommends the harmonisation of education across the ERA to support student, researcher and worker mobility, mutual recognition of qualifications and equal opportunities. Standards need to be developed for sharing training material and e-Infrastructure between institutions, as well as for student and teacher identification enabling access to e-Infrastructure facilities and resources use;
- The e-IRG recommends the formation of user community support infrastructures in each European country. It especially encourages the collaboration of several countries in order to increase the efficiency of such infrastructures.

18 ICEAGE: http://www.iceage-eu.org/library and EGEE: http://egee.lib.ed.ac.uk

3. *Grid and Cloud Computing*

Editor: Cristy Jane Burne <cristy.burne@gridtalk-project.eu>,
Contributors: Marc-Elian Bégin <meb@sixsq.com>, Neil Geddes <n.i.geddes@rl.ac.uk>, Fotis Karagiannis <fkara@aueb.gr>,
Reviewer: Ignacio Martin Llorente <llorente@dacya.ucm.es>, Peter Kunszt <peter.kunszt@cscs.ch>, Malcolm Atkinson <mpa@nesc.ac.uk>

In the move towards a sustainable pan-European e-Infrastructure, Europe must take care to choose and make best use of suitable IT technologies and service offerings. While academia has thus far approached distributed computing infrastructures using the Grid, useful for its federating ability and collaborative approach, the introduction of the new "cloud computing" paradigm demands a fresh assessment and attitude.

Clouds and grids[19] are currently developing in parallel and have strong potential to complement each other. An understanding of their relationship, and the ways in which it will continue to change, offers opportunity to aid the development and efficient use of existing academic and commercial e-Infrastructures.

3.1. Definition of policy areas

Europe's competitive global position is increasingly dependent upon the e-Infrastructures that power its dynamic knowledge-based economy. The shape these infrastructures take is influenced by a complex cocktail of technical and social possibility and opportunity, as well as commercial interests, political agendas and a range of vested interests. This section attempts to address the relation of the industry-driven area of cloud computing with the research-driven Grid-enabled e-Infrastructure, proposing policy actions for the enrichment and evolution of the latter.

Many research endeavours currently rely on access to and drive the development of grid-like e-Infrastructures that support regional, national and international collaborations through enabling the federation of their resources. Development of these infrastructures has been steady, but often faced with ongoing challenges relating to interoperation, standardisation and ease-of-use. In the wake of the relatively simplistic cloud solutions, which focus not on federation but on efficient resource provision, there is significant room for the research community to learn from and even adopt the technologies and service offerings clouds have introduced. Specifically, clouds could help in overcoming many barriers to grid adoption. In general, the mixture of the best of these technologies and services can create a new breed for the benefit of both researchers and plain users, for e-Science and e-Business respectively.

3.2. Goal of the policy analysis

This section aims to clarify understanding of the current situation, as it relates to grid and cloud services, for Europe's major e-Infrastructure players, including funding agencies, e-IRG and EGI members. Specifically, it aims to provide an assessment of the commercial understanding of the research community's e-Infrastructure requirements, as well as recommendations for enhancing the same. Equally, it aims to determine the best way forward for existing research and business e-Infrastructures, which need to focus on improving automation (to increase agility, while reducing operational costs), simplicity (both for users and providers) and diversity (freedom of choice provided by commercial and academic resources).

Which of the bespoke services used by today's research community might become commodities that can be efficiently procured from generic service providers? What action is needed to facilitate this? Which of the research community's needs are likely to remain qualitatively or quantitatively beyond the scope of commercial cloud providers? How can grid and cloud services complement and learn from each other? This section aims to address these questions.

3.3. Context

Grid computing is a relatively established form of distributed computing, adopted early by the high energy physics community and increasingly applied to other disciplines. In contrast, cloud computing has experienced a recent surge and is a service now offered by a growing number of IT companies. Grids and clouds share two main attributes: 1) they provide access to remote resources; and 2) they provide a non-trivial quality of service. Both grids and clouds have adopted the concept of IT "as a service," with grid more likely to offer to the end-user access to shared resources without direct end-user payments, while clouds currently have a pay-as-you-go approach. Thus Grids and clouds are at present very different in terms of a "business model": The Grid is more focused on enabling federation of distributed, heterogeneous resources, enabling joint usage and sharing of data by virtual organizations formed by user communities crossing national and/or administrative boundaries, Cloud computing focuses more on on-demand provision of IT infrastructure, initially targeting less high-end businesses.

19 Definitions are available in the Terminology section

Further, grid computing tends more towards the use of open standards, whilst clouds rely more often on proprietary interfaces.

The unique and complex requirements of the research and academic community have thus-far justified the existence of tailored and publicly funded grid-style e-Infrastructures. If cloud services offered by industrial players such as Google or Amazon go near to satisfying the advanced requirements of large scientific communities, then such grid computing e-Infrastructures could have a hard time justifying their existence. Grids must thus offer services that are one step ahead of those offered by industry, catering to specific demands that cannot be met with the same level of efficiency using the generic resource provision of clouds.

The relative cost of provisioning computing and storage resources via (outsourced) cloud services or (insourced) grid infrastructures has been compared in a number of research papers[20], leading to the conclusion that commercial clouds are "probably not yet" able to support the complex use scenarios, including flexible access control and support for collaborative access, delegation and auditing that the research community demands. Equally, questions have been raised about the security aspects of commercial offerings, which are currently unclear.

There has been an increasing number of "proof-of-concept" developments investigating ways of integrating cloud-based, on-demand services with existing grid-enabled applications across a number of disciplines, from nuclear physics to gene expression. Response from these projects regarding the potential benefits of integrations has been almost exclusively positive, but few have made the transition into everyday production[21]. Potential benefits include advantages both for end users (reducing the length and cost of application porting), and resource providers (reducing the operational cost of running grid sites).

As long as proprietary developments are seen to provide commercial advantages, there is little evidence that large providers will seek to standardise their offerings. There are, however, a number of emerging open source implementations of the underlying commercial systems from Google and Amazon. These may help to drive some standard application interfaces in the future. However, the success of cloud computing shows that carefully selection of underlying technologies can delay the need for standardisation and enable de-facto standards, promoted by strong support communities, as is the case, for example, with virtualisation.

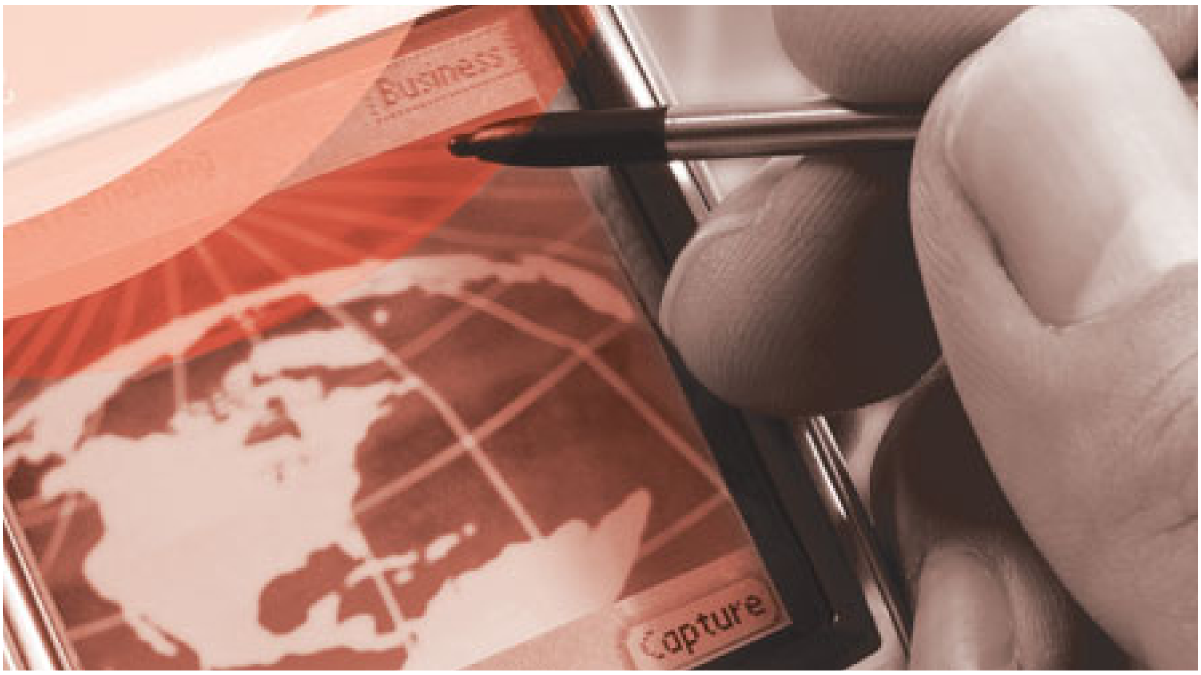

20 http://www.symmetrymagazine.org/breaking/2008/05/23/are-commercial-computing-clouds-ready-for-high-energy-physics/

21 http://www.isgtw.org/?pid=1001735

Recent developments in cloud computing have also given rise to significant changes in the scale of resources available as well as their availability or other quality of service characteristics, and the simplicity with which they can be accessed. Sustained by significant commercial investment, continued expansion of cloud resources is expected. Combined with growing competition between current providers, this should also help reduce the cost of access. This cost needs to be balanced, however, against the transitory and proprietary nature of the services provided. Other issues that would need to be addressed and clarified soon are ones dealing with trust in the provider or the country hosting the resources (since storing any kind of "sensitive" data remotely introduces potential risks), combined with legal issues, that might come up in case of violating foreign law.

3.4. Proposed approach

Commercial offerings do not yet fulfil the specific and demanding IT requirements of the research community, and any immediate cost savings are not thought to be sufficient to drive significant reengineering of these requirements. There is clearly a lot of interest in exploiting the growing number of cloud offerings, and support should be given to exploring the use of the emerging services and their integration into existing research infrastructures. Equally, cloud offerings should take advantage of the research and development conducted by the grid community, for example in the adoption of standard grid interfaces and functionality for federation.

An important ingredient in promoting a more efficient e-Infrastructure is to also advance the under-lying middleware, in a way to depart from single solutions. Simpler solutions, composed of an aggregation of focused simple services, are better than a single solution which, even if open source, restricts freedom of choice. The intrinsic complexity of the distributed computational context and diverse user requirements, already make middleware quite complex, something that needs to be hidden from the users.

Ultimately individual projects and agencies will make their own judgments on the relative cost and merit of cloud-based services as compared to grid-based provisioning. However, the long-term interests of the research community would clearly be served by access to and use of a mixture of grid and cloud-based services and related technologies. Such a development would:

- encourage further take-up of services as users gain confidence that their own investments will have a long-term future;
- avoid "lock-in" to particular proprietary solutions;
- allow new providers to offer similar services.

Alongside this approach, there is an ongoing need to promote the benefits of data center automation, with the acknowledged goal of increased efficiency through fabric management, system monitoring, resource virtualisation and development of simpler interfaces.

3.5. Recommendations

- The e-IRG acknowledges the apparent rapid growth of commercial Cloud-like services, providing on-demand virtual computing resources, data storage and software services. The e-IRG recommends that major e-Infrastructure initiatives investigate the integration of commercial and non-commercial infrastructure services and of Grid and cloud-like technologies especially for achieving the provision of on-demand virtual computing and storage resources into existing e-Infrastructures;
- e-IRG recommends that the EC and MS support should not be limited to a single distributed computing technology and infrastructures, promoting an open approach, when aiming to set up sustainable pan-European e-Infrastructure organisations within Grid, Cloud and High Performance Computing.

4. Security: a holistic approach

Editors: Fiorenzo Scaroni <fiorenzo.scaroni@switch.ch>, Christoph Graf <christoph.graf@switch.ch>
Contributors: Christoph Witzig <christoph.witzig@switch.ch>
Reviewers: Neil Geddes <n.i.geddes@rl.ac.uk>, Christos Kanellopoulos <skanct@grid.auth.gr>, Diego Lopez <diego.lopez@rediris.es>, Milan Sova <sova@cesnet.cz>, Daniel Kouril <kouril@ics.muni.cz>

Security matters. Users and service providers will adopt new e-Infrastructure only if it is perceived to be fit for the purpose and supports appropriate trust and security levels. Offering an adequate level of security for a wide range of uses is a hard problem for many reasons:

- Security is hard to describe and measure: there exists no commonly accepted language and metric for all aspects of security and it is often judged on subjective grounds only;
- Security has different meanings depending on the domains and the scope;
- Security is costly: those willing to accept a lower level of security are generally not willing to pay more simply because others require a higher level;
- Security can be cumbersome: a higher level of security often makes it harder to use a piece of e-Infrastructure;
- Security is a moving target: what is secure enough today, might not stay secure enough tomorrow. Sometimes, such change occurs almost instantly: new discoveries might instantly render important security elements useless;
- Security has many stakeholders: an attacker needs to find only one particular method to successfully attack a service. The defender, on the other hand, needs to protect against all possible attacks to the service, covering the entire supply chain and involved parties.

4.1. Definition of policy areas

While it is important that security is addressed in its entirety, modularisation can help to make security more manageable. This requires harmonisation of the policies at different levels of abstraction. We begin with a coarse but essential taxonomy. Elements with reasonably well-understood and limited policy impact can be addressed separately. This method should be used to address the following areas individually:

- Reactive security: Protection of the network infrastructure and overlaid services based on the willingness to be able to react to any event, which might affect the operation of the network, within a short space of time (incident handling);
- Proactive security: Promotion of awareness and cooperation through education and through offering a platform for sharing of information, experiences and knowledge between the responsible people within the community in order to identify areas of common interest for further cooperation (threat and vulnerability management policy);
- Access management: "The process responsible for allowing users to make use of IT services, data, or other assets. Access management helps to protect the confidentiality, integrity and availability of assets by ensuring that only authorized users are able to access or modify the assets. Access management is sometimes referred to as Rights Management or Identity Management". (ITIL definition[22]).

4.2. Goal of the policy analysis

This chapter emphasizes the role of security as a central plank of e-Infrastructure. Security is a fundamental and indispensable component that should become an intrinsic, positive factor of e-Infrastructure. The development of policies and implementation has, to date, taken place independently in different domains. These domains align, to a large extent, with several flagship projects financed by the European Commission. We propose a framework as a working base for the implementation of security services.

4.3. Context

The understanding of security varies broadly depending on the domain of the e-Infrastructure. Each domain has different aims. By way of example we can list some relevant aspects of four representative fields of applicability:

- National Research and Education Networks (NRENs) are based on a national scope. Reactive and proactive security has a tradition of collaboration built over many years, which led to a close collaboration between the Computing Security Incident Response Teams (CSIRTs). Access management is being addressed successfully at national level by means of federated Authentication and Authorisation Infrastructures, which have privacy preservation as one of their main objectives. The international interoperability is now addressed through interfederated access management;

- Grid providers support large project-oriented communities with no institutional or national boundaries. The Grid community, being early adopters of the concept of large-scale computing across administrative domains adopted a security framework, which is largely incompatible with emerging AA-solutions. Today, interoperability issues arise at technical as well as organisational levels where the chosen security model is at odds with institutional or organisational models;

22 http://www.knowledgetransfer.net/dictionary/ITIL/en/Access_Management.htm

- Unique specialised resource providers, e.g. supercomputers, large accelerators, national libraries. They adhere to a closed security model, having little interaction with other resource centres and communities;

- Data aggregators, not yet a clearly contoured domain, spans small and closed to very arge open user communities. Today the data issues are not treated as a generic domain but managed by application specific communities.

The sectoral security policies and their implementations represent achievements, which can accommodate the local needs of different communities. However, the growing interconnectedness of research and education as well as of industry and services on a European level, and worldwide, induces an overlap of underlying security services.

First steps to bridge the different e-Infrastructure stakeholders (service providers, user communities and policy makers) are being taken at a number of levels, for example:

- NREN and Grids Workshops organised by TERENA;
- The REFEDS[23] group, gathering stakeholders from several communities worldwide;
- Developments to make interoperable federated identity management systems (such as Shibboleth) with Grids (e.g. in EGEE, Globus, D-Grid and others);
- IGTF as the forum to coordinate Certification Authorities in Europe, Asia and the Americas in the academic sector;
- TF-CSIRT to network CSIRT teams within Europe.

The aim of these and similar initiatives is to promote intersectoral collaboration. They can be viewed as important starting points in the development of common policies or understood interfaces between persistent co-evolving interfaces.

4.4. Proposed approach

The main challenge over the next few years will be:

- The definition of accepted security specifications and metrics, which will serve as common ground and a framework for all activities;
- The collaboration of different key players in the field of security, spanning both policy makers and e-Infrastructure service providers implementing security services;
- Promotion and focus of **federated** approaches across scientific communities with key national elements acting as glue between autonomous local centres and communities on one hand and European and global efforts on the other hand.

As an example, in the area of trust management, it appears that not all identity or service providers provide the same "level of assurance" and it may be challenging to find a peer that provides the requested level. Progress has been made in the CA domain, where the IGTF coordinates the accreditation process, but similar mechanisms should also be made available for the identity (inter)federations as well.

Policy aspects should address:

- Reactive Security: incident management policy;
- Proactive Security: threat and vulnerability management policy;
- Access Management: specifically to interoperate with established and emerging AAA Infrastructures.

23 http://www.terena.org/activities/refeds/

A good security framework should, where possible, be independent of particular domains. It must serve needs efficiently on different levels of abstraction:

- Project and virtual organisation;
- Organizations as individual research and education institutes;
- National and political boundaries;
- International level, where today the most activities in research and education as well as in industry and services take place.

A coordination and orchestration of security related activities in these areas across the flagship projects financed by the European Commission and other bodies (e.g. EGEE-EGI, DEISA-PRACE, data repositories projects, OMII, IGTF, OGF) and other relevant stakeholders (e.g. NRENs and TERENA) is essential if a coherent security framework is to be achieved.

4.5. Recommendations

- The e-IRG strongly encourages the harmonisation of approaches in Access Management between the NREN and Grid e-Infrastructure providers;
- The e-IRG recommends evaluating whether the well-established CSIRT collaboration platform in the NREN-community could possibly accommodate the needs of other e-Infrastructure components before development of domain specific structures for analogous tasks;
- The e-IRG recommends the promotion of a coherent framework that will act as a working base for the implementation of unified security services. In particular, the full extent of cross-disciplinary synergies of a coherent model for data security should be studied and exploited.

5. Service-centric e-Infrastructures through virtualisation

Editor: Radek Januszewski <radekj@man.poznan.pl>,
Contributors: Mauro Campanella <Mauro.Campanella@garr.it>, Fotis Karagiannis <fkara@aueb.gr>,
Reviewers: Ludek Matyska <<ludek@ics.muni.cz>, Jesus Marco <marco@ifca.unican.es>, Mauro Campanella <mauro.campanella@garr.it>

Currently the growth of costs of ownership of a large number of independent physical systems outpaces the benefits of the growing performance of the modern hardware installed in pan-European e-Infrastructures. Increasing energy and natural environment costs are forcing us to rethink the way the available resources are used. The traditional way of engineering e-Infrastructures, relying on dedicating the hardware for a single purpose and a single user group, is no longer valid due to security and political reasons. Sharing resources may lead to significant reductions in the cost of ownership and in an increased efficiency of its utilization by sharing it.

At the same time, Internet has become an increasingly important element of the global economy. This fact coupled with the trend of using shared infrastructure (Grids, Cloud Computing) and the growing importance of confidentiality and security of processed information makes the case for an isolation of customer-dedicated parts of the infrastructure.

So there are two opposing trends – to share and at the same time to be secure and confidential; so the question is how do we achieve the two goals at the same time? The answer may be the technology of virtualisation. The physical resources host virtual environments customised to exactly match user needs, dynamically utilizing only required resources and allowing other users to access and share the unused ones.

Virtualisation technologies, both in computers and in networks, have been available for a long time. In the last years virtualisation is receiving particular interest and hardware manufacturers are now embedding in products features to assist virtualisation. Virtualisation can be for

simplicity defined here as the use of technologies to create on the same set of physical resources more than one logical environment perceived by the resource consumer as isolated and performing as if the consumer was the only user.

5.1. Definition of policy areas

Virtualisation aims at and allows a more efficient and flexible usage of resources and it can significantly contribute to the way future e-Infrastructures will be utilized and shared. Currently the term is understood differently in the context of different technologies such as storage, processing, networking etc. To achieve a fully innovative environment, research effort is needed to define common standards to build bridges across technologies and facilitate interoperability.

Virtualisation technology, especially if it is used in a coordinated way in the areas of computing and networking, can be essential for solving some of the problems of the contemporary e-Infrastructures and make resource sharing in the European projects much effective. It allows for rapid deployment of new environments and helps to automate and reduce the management and administration processes.

Virtualisation may also offer greater reliability and independence of the functionality (e.g. a web server) from its physical location using logical migration and distributed storage. The increasing prices of electricity used for powering and cooling e-Infrastructure are becoming a larger part of the total cost of ownership and maintenance of the computing infrastructure. Dynamically adjusting the number of powered nodes to match the current requirements and levelling the load of each node to maximize utilization of powered infrastructure can lead to significant advantages both in terms of *reduction of the cost of ownership and in the ecological benefits* of using the infrastructure. The security, confidentiality and integrity of computing and the results in the competitive world of science are becoming the key features of the future pan-European shared computational e-Infrastructures and the best way to provide these qualities is maximum isolation of the applications utilizing the distributed infrastructure.

Another major aspect of the technology is the possibility of increasing the diversity of the used software. It is not uncommon to have an application designed/developed in a specific operating system. Up to now one had to provision a dedicated computing resource to run specific applications. The need for using certain applications that is crucial for some group of users may entail additional purchases of hardware as the various environments could not share the hardware resources, even if its performance was enough to host and run all the applications at the same time. With virtualisation the different operating systems may share the same hardware in a secure and efficient way. This is increasing competitiveness between the system providers, as the user is no longer bound to install and use software delivered by one provider on his computing resources. It is possible to prepare various dedicated environments for different applications sharing the underlying hardware and thus reducing costs. Virtualisation may then *offer a uniform interface to basic computing and networking resources*, adapting easily to novel uses and allowing e-Infrastructure modification and upgrade (e.g. newer processors, storage, networking) not noticed by the user.

5.2. Goal of the policy analysis

The main objective of this section is to assess the applicability of virtualisation in e-Infrastructures and present policy benefits and implications. The (virtualisation) topic can be relevant to e-Infrastructure projects related to a variety of layers such as networking, computing, storage, presentation and applications. An effort to identify the applicability of virtualisation to the above e-Infrastructures layers is attempted in this section. More detailed objectives include:

- To introduce new, state of the art techniques for resource management;
- To promote environment-friendly solutions as a standard in the e-Infrastructure;
- To make sharing of the resources more secure;
- To promote research on the next generation environments.

5.3. Context

Currently this technology is used mainly by industry because economic benefits are the most obvious ones and the most easy to achieve. There are, however, many examples of successful deployment of virtualisation in supercomputing centres resulting in better resource utilization, increased flexibility of the maintenance and introduction of additional unique features.

The main actors in this area are companies such as VMware, Microsoft, SUN, IBM offering a basic environment free of charge and featuring rich-commercial solutions, but there are also open-source projects such as Xen and LVM offering similar functionality. It is worth noting that because of the dominant role of Linux based OS in the world of scientific computing the solutions based on the open-source projects are more and more commonly deployed and the tools are maturing rapidly.

There are projects that are using virtualisation for Future Internet research like FEDERICA (www.fp7-federica.eu), the GENI initiative in the US (www.geni.net) and AKARI & JGN2plus in Japan. The interest in Future Internet research is increasing from multiple stakeholders such as the EC and individual countries. In parallel, infrastructures are being built (e.g. FIRE initiative, G-LAB in Germany). Moreover future internet research involves all layers from mobiles to applications. The Open Grid Forum - Grid and Virtualisation Working Group (http://www.ogf.org) also deals with this area as well as the Distributed Management Task Force (http://www.dmtf.org). The FEDERICA test-bed is built on top of the existing NRENs and GEANT networks linked in a federated environment. FEDERICA will not deploy all the possible technologies in its own test-bed. It plans to interconnect/federate with other infrastructures to overcome this limitation.

Traditional client-server computing is starting to lose ground as a new paradigm emerges – the Cloud Computing paradigm. Cloud Computing allows data centers to offer services transparently through the Internet by exploiting their computing and storage fabric of resources. Services are delivered on demand over the web, reducing software complexity and costs, expediting time-to-market, improving reliability and enhancing accessibility of consumers to government and business services. In the Cloud Computing paradigm virtualisation technology is essential because it deals with multiple layers, such as virtual networking, virtual storage, virtual presentation and virtual applications.

This concept of using resource virtualisation allows for shifting from a hardware-centric to a service-centric view of computing where all the underlying hardware, be it storage, computing or networking, is perceived as a pool of resources that one can allocate dynamically to exactly meet the current need for services, reducing energy consumption and costs.

5.4. Proposed approach

An effort should be made by the major e-Infrastructure projects dealing with the identified layers, i.e. networking, computing, storage, presentation and applications to assess the implications of and benefits from using virtualisation technologies. The FEDERICA EC co-funded project will create an agnostic Europe-wide infrastructure to support Future Internet research experiments and advance knowledge on management of virtualized resources (network and computing elements). Similar efforts could be investigated in the capacity computing projects, i.e. EGEE-EGI. For the capability computing projects, i.e. DEISA-PRACE virtualisation does not seem to be in-line with perceptions, as they offer a big number of resources and this technology does not appear to be useful. However, virtualisation techniques could be used in other layers such as storage, presentations or applications.

The RESERVOIR EC-funded project aims at supporting the emergence of Service-Oriented Computing as a new computing paradigm. Virtualisation plays a significant role in this task because this technology has proven to be useful in overcoming some barriers to commercial adoption of grid technology. As highlighted above, the cloud computing model extensively uses virtualisation techniques and thus is in the core of this topic.

5.5. Recommendations

- The e-IRG notes the emerging use of virtualisation in ICT service provision, whereby physical resources can be shared by users in a manner which appears to support each user independently, optimizing resource utilization, reliability, energy efficiency and maintenance costs. The e-IRG recommends the investigation of virtualisation in key e-Infrastructure projects;
- The e-IRG recommends that further research on virtualisation concepts is supported, including development of open standards for integration of tools from different vendors and academia in order to support emergence of a competitive marketplace in this domain.

6. *Remote instrumentation*

Editor: Norbert Meyer meyer@man.poznan.pl, Contributors: Marcin Lawenda <lawenda@man.poznan.pl>,
Reviewer: Jesus Marco <marco@ifca.unican.es>, Costas Cotsokalis <Constantinos.Kotsokalis@udo.edu>

The term instrumentation is widely used for laboratory equipment which is necessary for daily use in experimental sciences, e.g. chemistry, biochemistry, physics, astronomy, but also by industry solutions. Instrumentation is, however, not only limited to laboratory equipment but may also cover various sensors which can be remotely controlled and monitored. Implementation of Remote Instrumentation Services (RIS) allows us to operate the instrumentation remotely by means of the location of the end user and the equipment. The necessity of using interactively advanced, unique and expensive equipment infrastructure, which is often locally unavailable, as well as the advantages of broad international cooperation are the key issues for the success of a great number of scientific disciplines. Remote assisted use of scientific equipment will substantially reduce the human and financial cost, as it provides the possibility of sharing the results in real time with other researchers at different geographic locations, making a better use of senior researchers' time.

Thus the development and spread of remote instrumentation techniques and technologies that allow virtualized, remote, and shared access to such infrastructure opens up new opportunities for scientific communities. Grid technologies facilitate the sharing of these resources across a distributed computing environment, and its integration as part of an e-Infrastructure. In particular, the Grid handles issues of authentication, authorization, resource description and location, data transfer and resource accounting. Moreover Grid technologies can be used to integrate operations with computing resources where complex models of the experimental facility can be executed.

6.1. Definition of policy areas

As stated before, we aim to consider Instrumentation as a part of European and Worldwide e-Infrastructure, which is embedded and strictly coupled with other components, starting with the network and including computing systems and data repositories.

The main policy issues relevant to remote instrumentation are the following:

- Enabling fair access to the instrument by remote users, including the preparation of generic models for service level agreement. This covers the definition of a fair and sustainable sharing scheme, where using remote instruments implies on one hand sharing own instruments, but on the same time under-privileged institutes that do not possess relevant equipment are not left out;
- Promoting the use of an integrated e-Infrastructure which encompasses various layers of the e-Infrastructure, such as networking, grid computing resources (computing, storage), grid middleware and the instrument itself. This point could be handled through the integration of an "Instrument Element" (IE) and its associated services in the grid middleware, in a similar way to the Computing (CE) or Storage Element (SE) services;
- Empowering the daily work of less privileged scientific communities (by connecting the remote instrument to the e-Infrastructure).

Other policy topics like the use of acquired knowledge and intellectual property rights, including the consideration of an Open Science framework, will have to be considered in the future.

6.2. Goal of the policy analysis

The principal aim of the "Remote Instrumentation" section is to provide an overview and a proposed approach concerning the integration of scientific instrumentation in the global e-Infrastructure framework, and how to make it available to research communities. This will result also in the reduction of the number of isolated islands of research infrastructure.
This potential access to all resources will also have an added value for the selected scientific communities, through the use of enhanced services for new needs:

This potential access to all resources will also have an added value for the selected scientific communities, through the use of enhanced services for new needs:

- Using the new generation of network services provided by the GEANT2 infrastructure, like Quality of Service and features offered by IPv6;
- Building a production quality distributed computational and visualization infrastructure, integrated with the distributed instrumentation resources of the selected scientific area;
- Providing and supporting tools enhancing the added-value to the e-Infrastructure (like interactivity, parallelization, collaborative tools, dynamic measurement scenarios - workflows, low latency).

6.3. Context

The new technology focuses on the creation of **new possibilities** for Remote-Instrumentation-based cooperation among scientists from different countries, working in the same or closely related scientific domains. The possibility of remote access will allow **better utilisation** of the instruments, therefore making resource usage more efficient. The owners will also have a chance to **cooperate** on their research with specialists from many different remote locations. The most obvious advantage of this new approach is that one can **exploit the storage and processing** capabilities that the classical data/computing ecosystem offers. By representing the instrument as a service and integrating it with other services through well-understood protocols, it becomes (conceptually, but also technically) straightforward to directly store experimental data to arbitrary locations world-wide, replicate them in multiple locations, and perform post-processing that might previously take days or months in only a fraction of the time. The massive Grid capabilities, surpassing any kind of supercomputer performance for specific categories of scientific problems, are a perfect match for experimental and applied science. Instrumentation as a service, an analogy to software as a service, allows composing atomic (in the software sense) experimental actions into measurement chains or long-standing experimental processes, irrelevant of the location of the cooperating instruments. Workflow execution mechanisms bringing together multiple different instruments, possibly geographically dispersed, can raise experimental science to new levels.

Current practices and achievements can be found in some of the well-known related projects such as the WLCG (HEP computing) and e-VLBI (experiments in radio astronomy in near real-time using GEANT2 and grids for distributed correlation) projects.

Several other communities require integration of instruments (sensor networks, equipments) with the IT environment – this seems to be a significant benefit for them. There has been significant experience for almost 7 years in all strategic areas that remote instrumentation will use, such as interactivity in Grids and distributed environments (e.g. Int.EU.Grid[24], BalticGrid[25], and CrossGRID[26] projects), visualization (e.g. Int.EU.Grid, CrossGrid, and g-Eclipse[27]), collaborative tools (e.g. GRIDCC[28]), steering, monitoring and accessing remote instrumentation (e.g., DORII, RINGrid[29], GRIDCC, VLAB[30], and CRIMSON[31]), building software frameworks for fast prototyping, easy and fast adaptation (e.g., DORII, g-Eclipse, and Int.EU.Grid), workflows (VLAB), networking, quality of service and bandwidth on demand (GN2 project[32]). Another effort made in the U.S. is CIMA – the Common Instrument Middleware Architecture, which aims at Grid-enabling a wide range of scientific instruments and sensors to enable easy access to and sharing and storage of data produced by these instruments and sensors.

6.4. Proposed approach

Developing a dedicated system for each piece of equipment would be very hard and inefficient. A better solution seems to be the gathering of similar instruments and trying to design general solutions for them. In order to get some degree of generality, a common use-case model for remote instrumentation must be defined. The first approach was made during the RINGrid project, a design study research in cooperation with various scientific communities and the participants of the Open Grid Forum RISGE-RG (Remote Instrumentation Services in Grid Environment) research group.

The next step towards achieving a general concept in remote instrumentation should be an analysis of standards.

Thus the development and spreading of remote instrumentation techniques and technologies that allow virtualized, remote and shared access to such infrastructure opens up new opportunities for scientific communities.

24 http://www.i2g.eu/

25 http://www.balticgrid.eu/

26 http://www.eu-crossgrid.org/

27 http://www.geclipse.eu/

28 http://www.gridcc.org

29 http://www.ringrid.eu

30 http://vlab.psnc.pl/

31 http://crimsonproject.net/

32 http://www.geant2.net/

Remote assisted use will substantially **reduce the human and financial cost** of using remote instruments, as it provides the possibility of sharing the results live with other researchers at their home institution or even at other institutions, making better use of senior researchers' time. Another alternative form of using information and communication technologies to make more effective use of expensive remote instruments is to **archive scientific data** resulting from observations in places not necessarily close to where they are generated, and ensuring efficient access to these archives.

Still another aim is to eliminate the current isolated islands of research infrastructure, by **integrating them into one big e-Infrastructure**, giving potential access to all resources (e.g. distributed computing, storage, visualisation infrastructure) and providing added value for the selected scientific communities, in terms of new services currently existing, but adopted for new needs.

Detailed objectives are the following:

- To use the new generation services provided by the GEANT2 networking infrastructure, i.e. quality of service, IPv6;
- To build a production quality distributed computational and visualisation infrastructure, integrated with the distributed instrumentation environments of the selected scientific area;
- To provide and support the added values of e-Infrastructure (like interactivity, parallelization, collaborative tools, dynamic measurement scenarios - workflows, low latency) that would be used in the integrated environment of scientific and engineering instrumentation, networking, visualisation and computational infrastructures;
- To **virtualise** instruments to make the approach more general as it is done for virtual organisations in grid environments and virtual resources in distributed computing.

The current activities performed in reference to this subject cover radio-astronomy in EXPReS project (http://www.expres-eu.org/), earthquake, environment and experimental science communities with reference installations mentioned by ESFRI – activities in DORII (http://www.dorii.eu). Standardisation approaches are under investigation by the RISGE research group in OGF and INGRID conference (Instrumenting the Grid, http://www.ingrid.cnit.it/)

The definition of a common use-case model for remote instrumentation and related standardisation will promote fair access to resources, including access of less privileged users, as well as promote the use of a fully integrated e-Infrastructure encompassing instruments (besides network, computing, visualisation and data infrastructures).

6.5. Recommendations

- The e-IRG recognises the benefits of remote instrumentation, especially for the large and costly research infrastructures;
- The e-IRG recommends the development of mechanisms enabling fair remote access to state-of-the-art equipment, including the preparation of a sustainable sharing scheme, both for public and private research, and research on standards for integrating remote instrumentation in the current e-Infrastructure (grids, advanced computing and data repositories);
- The e-IRG invites the EC and member states to strengthen the support for research on remote instrumentation. This should cover technological, economic, policy and security aspects.

7. Towards sustainability of the computing-related e-Infrastructure

Editor: Fotis Karagiannis <fkara@aueb.gr>,
Contributors: Mirco Mazzucato <mirco.mazzucato@pd.infn.it>, Peter Kunszt <peter.kunszt@cscs.ch>, Thomas Eickermann <th.eickermann@fz-juelich.de>
Reviewers: Neil Geddes <n.i.geddes@rl.ac.uk>, Sergi Girona <sergi.girona@bsc.es>, Rosette Vandenbroucke <vandenbroucke@ helios.iihe.ac.be>, Peter Kunszt <peter.kunszt@cscs.ch>, Gabrielle Von Voigt <vonvoigt@rrzn.uni-hannover.de>

One of the key issues to broad cross-disciplinary adoption and use of the e-Infrastructure by the user communities is to trust that the e-Infrastructure services are reliably available and steady for the estimated duration of their current and planned collaborative projects. To achieve this in an e-Infrastructure provided by a cooperative activity invoking actors across national borders requires a common perception and implementation of cost-sharing models and related policies. At the moment Europe is aiming at adding capability (HPC or supercomputing-like) and capacity (Grid-like) computing services to its existing sustainable e-Infrastructure provided by the GÉANT and the NREN networks. This section deals mainly with the policy issues related to the interworking of the capability and capacity computing. Smooth interworking and interoperation of these two components is important for the sustainability of the computing-related e-Infrastructure.

7.1. Definition of policy areas

The main topic that needs to be tackled in the near future is the transition approach towards developing a sustainable computing layer on top of the research networking layer. As identified above, this computing layer deals with both the capacity and capability computing, which need to smoothly and seamlessly interact inside the computing ecosystem. Currently there are two main pan-European projects working towards this objective:

- The European Grid Initiative Design Study project – EGI_DS (www.eu-egi-eu), a design study aiming at preparing for a sustainable organisational and operational structure for Grid related services (capacity computing) in Europe;
- The Partnership for Advance Computing in Europe – PRACE (www.prace-project.eu), a preparatory project aiming at preparing a persistent petaflop-scale supercomputing service (capability computing) in Europe.

Both efforts will identify the legal, administrative and operational structures required for the persistent and sustainable running of the corresponding computing and data services. The main issue analysed in this section is the **interworking** strategy of the two components of the computing e-Infrastructure layer. Further policy issues can be identified for each of the two components such as the business model and funding schemes for this new modus operandi, basic commitments from the involved stakeholders (such as National Grid Initiatives / National Supercomputing Centres or Governments / Research Councils that need to commit funding and/or infrastructures) along with middleware related issues (such as selection policies of middleware distributions, creation of middleware consortia and interoperability criteria). Both projects have been preparing detailed proposals (deliverables and blueprints) detailing most of the above issues and calculating the required funding for the development of each of the components, along with a strategy to secure a consistent long term financial basis.

7.2. Goal of the policy analysis

The main goal of this section is to present the way of smooth interworking between the two identified computing components of e-Infrastructure (capability and capacity), focusing primarily on policy-makers, but also taking into account the service providers and user communities' views. In addition, an effort is made to reveal some side policy issues related to each of the two components' sustainability.

7.3. Context

There are basically two different views on the issue of the creation of the computing ecosystem in Europe, each one developed and promoted by the two components' communities. Hereafter follows a short analysis of the two views.

EGI_DS view:
The main foundations of EGI are the National Grid Initiatives (NGIs), which will provide on the national level the services for a seamless, shared and uniform access to a variety of computing resources, ranging from desktop computers and clusters, to sites operating supercomputers and all sorts of scientific archives. In other words, EGI_DS does not distinguish between the various types of resources but focuses on the integration and interoperability aspects. Besides NGIs, EGI will be composed of a central, mostly coordinating legal entity, called EGI Organisation (EGI.org or EGI.eu). EGI.eu will not provision or own any of the resources; these will be owned by the NGIs. The above are depicted in figure 7.1.

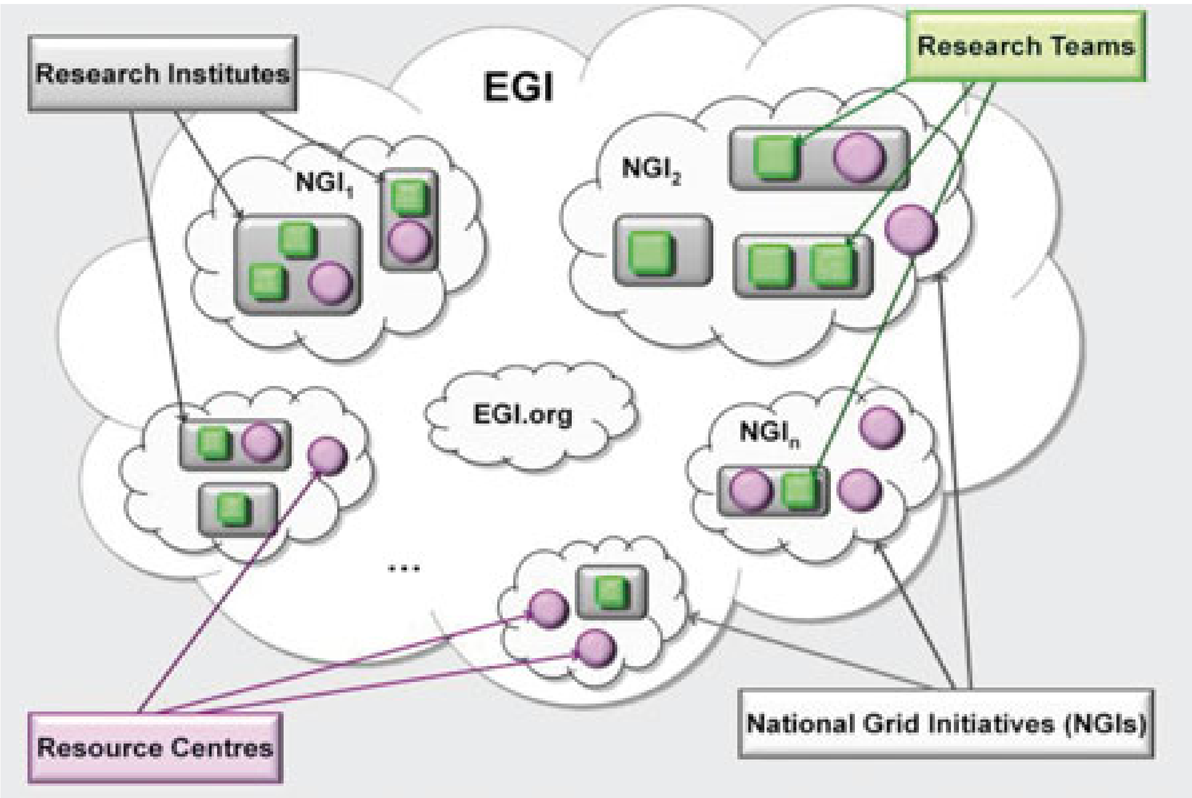

Figure 7.1 : EGI view of the computing ecosystem[33]

The EGI's added value is the ability of NGIs to share and integrate their heterogeneous resources that can be used by a variety of user communities through common Grid middleware interfaces and services. NGIs hence see themselves playing a central role at the national level, encompassing both the EGEE-type of resources and possibly local supercomputers not allocated to other stakeholders. The situation is similar for other current or future international facilities such as ESRF, LHC, ILL, ESS, ITER, where computing resources of such facilities are first integrated at the national level, usually through agreements between the country and the head institute for the commitment of resources over time. On the other hand PRACE is preparing a separate legal entity to provision (i.e. own), operate and manage 3-5 petaflop supercomputers that will be built around Europe. As a provider of a transnational access infrastructure, EGI wants to collaborate with other stakeholders and provide a standard, uniform access to all types of ICT resources. An idea that needs to be discussed is to agree on an EGI-PRACE peering or interoperation with the PRACE petaflop service. The EGI_DS project thus sees the NGIs and EGI.eu as main interworking actors.

Integration at the national level might have various effects in different countries. From current experience it appears that in countries with fewer resources and simpler structures integration is straightforward, while bigger countries might face serious problems in integrating the different stakeholders. In many cases the challenges stem from inherent funding policies e.g. different sources of funding (such as ministries) for the different components of the computing ecosystem and conflicting interests. Finally, more complicated structures might need to be put in place to run this integrated national computing environment.

PRACE view:
From the beginning, the PRACE project has adopted a tiered (i.e. layered) approach to the computing ecosystem (figure 7.2). It is a picture of resource ownership and available resource size,

33 Source: EGI Blueprint (http://www.eu-egi.eu/blueprint.pdf)

i.e. the tiers represent different levels of resources provision and coordination, as well as system performance. It is envisioned that the 3-5 petaflop systems owned by the PRACE legal entity build the top layer (Tier-0) of the performance pyramid and provide their services at European level. The national supercomputing systems are placed at the Tier-1, owned by the corresponding country or by large user communities. The rest of the capacity computing resources are owned or coordinated by the various NGIs and encompass larger cluster resources and end-user desktops at the two bottom layers. Besides operating and owning the Tier-0 systems, PRACE sees itself interacting with the Tier-1 resources, the national supercomputers. The Tier-1 resources currently integrated by DEISA, and the related services can also be used and integrated in the PRACE resources. PRACE takes no direct view on coordination of resources at the lower levels of the pyramid. Since the models and policies for resource allocation are typically significantly different between capability and capacity computing systems, PRACE does not envision an integration of all tiers or common governance, but rather a smooth interworking, allowing users to seamlessly access resources of different tiers.

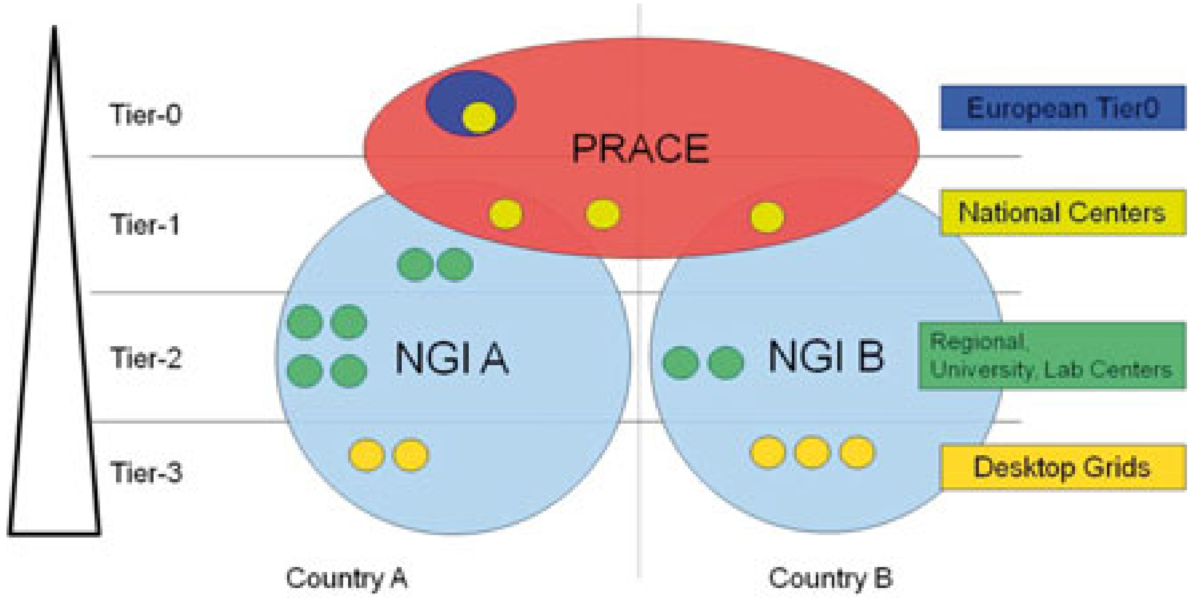

Figure 7.2: PRACE view of the computing ecosystem[34]

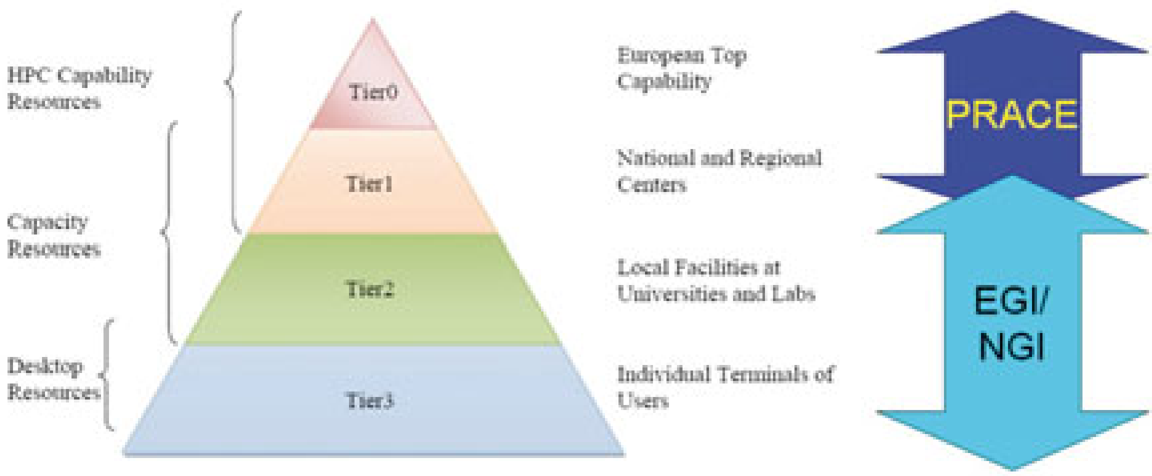

Figure 7.3: PRACE view of the computing ecosystem and organisational interfaces[35]

34 Source: PRACE Consortium
35 Source: PRACE Consortium

7.4. Proposed approach

Both EGI and PRACE work on regulating trans-national resources; In the case of PRACE under direct control, for the 3 to 5 petaflop systems and possibly a portion of the national supercomputers. It is expected that the PRACE resources are always over-allocated and there is no concept of resource sharing. Access to the resources is governed by a peer review based on the merit of the application (scientific relevance and technical feasibility) in using the HPC system. PRACE does not foresee addressing the issues with distributed storage management to the extent the EGI does.

In the case of EGI regulation of resources is done under delegated control, in a Grid model fashion. Access to the resources is mainly based on the need of massive transnational flows of data, access to data bases and collaboration services. EGI.eu aims to operate no resources itself, adopting the role of integrator and coordinator of resources, whether national (e.g. NGIs) or international (e.g. LHC, ITER). The EGI services offered to VOs for the management of distributed storage and data will not compete with any PRACE offer. As provider of an access infrastructure EGI can offer to PRACE to coordinate/ integrate the EU efforts aiming at providing standard uniform access to all types of ICT resources.

For the benefit of the user communities who need to see seamless services and interfaces, both communities need to work closely together in the future to define an acceptable solution enabling the two computing components to interact and interoperate in a complementary fashion.

Given the above analysis, EGI.eu could, therefore, play a key role in the interoperability efforts. The agreement for coordination between ARC, gLite and UNICORE Middleware Consortia aiming at full interoperability and a common European distribution called Unified Middleware Distribution (UMD) goes also in this direction.

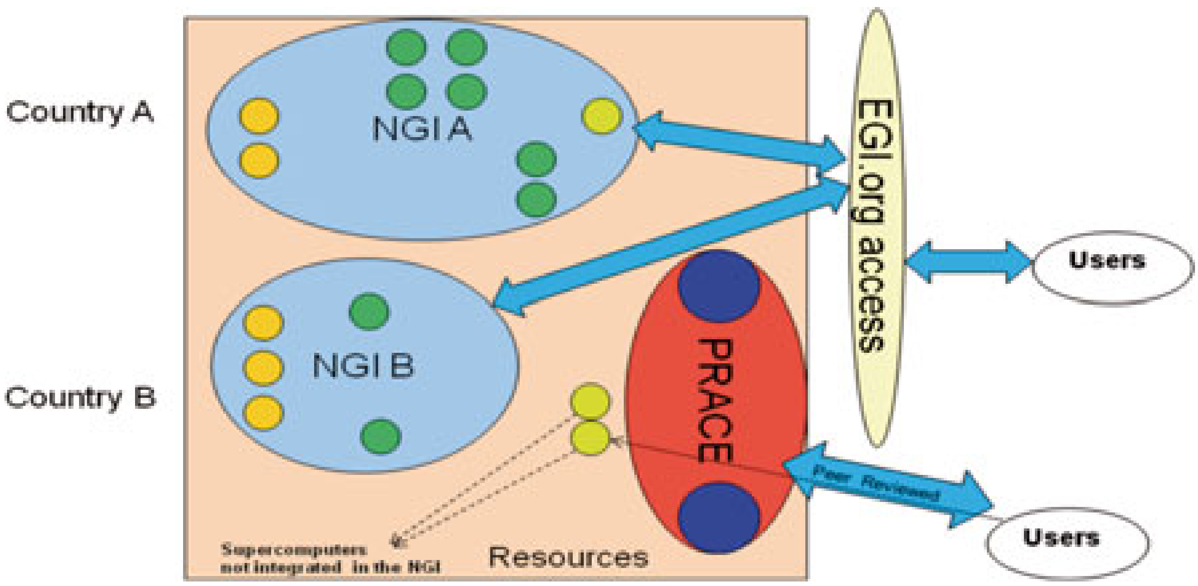

figure 7.4: A possible interaction scenario[36]

36 Initial source: PRACE consortium reworked by editor

7.5. Recommendations

- The e-IRG notes the importance of the steps undertaken by the EGI and PRACE initiatives to promote sustainability of the computing-related e-Infrastructure such as the development of policies, business models and funding schemes for the new required structures. e-IRG recommends that adequate levels of funding should be granted by the EC and Member States for the development of the new structures both on national and European level;
- The e-IRG recommends that major e-Infrastructures initiatives such as EGI and PRACE cooperate closely in order to define a complementary and interoperable environment for the benefit of European researchers. This environment should ensure that access to resources in Europe is granted through an open and transparent process, based on international standards and interoperable middleware;
- The e-IRG recommends the funding of activities that help national user communities to cooperate with corresponding user communities in other countries in order to foster the European research activities in using the e-infrastructure.

Editorial Responsibilities

Foreword: Editor Leif Laaksonen <leif.laaksonen@csc.fi>

Introduction: Editors Matti Heikkurinen <matti@emergence-tech.co.uk>, Fotis Karagiannis<fkara@aueb.gr> Review: Elizabeth Vander Meer <evmeer@nesc.ac.uk>

Terminology: Editors Radek Januszewski <radekj@man.poznan.pl>, Fotis Karagiannis, Elizabeth Vander Meer, Marie Christine Sawley <sawley@cern.ch>, Torbjorn Wiberg <torbjorn.wiberg@adm.umu.se>

Executive Summary: Editor Matti Heikkurinen, Damien Lecarpentier <Damien.Lecarpentier@csc.fi>, Fotis Karagiannis Review: Elizabeth Vander Meer

Overall Reviewers: Malcolm Atkinson <mpa@nesc.ac.uk>, (Sections 1, 2), Neil Geddes <n.i.geddes@rl.ac.uk>, (Sections 1, 2, 3, 5), Matti Heikkurinen, Kees Neggers <Kees.Neggers@SURFnet.nl>, Marie Christine Sawley, Elizabeth Vander Meer, Damien Lecarpentier

Section 1, Global Collaboration
Editors: Marie-Christine Sawley, Ana Bela Dias <dias@nwo.nl>, Contributors: Bill Feiereisen <bill@feiereisen.net>, Fotis Karagiannis Reviewers: Kees Neggers

Section 2, Education and Training in the use of e-Infrastructures
Editor: Malcolm Atkinson, Contributors: David Fergusson <dfmac@nesc.ac.uk>, Elizabeth Vander Meer, Jo Newman <jo@nesc.ac.uk>, Reviewers: Neil Geddes, Rosa Badia <rosa.m.badia@bsc.es>

Section 3, Grid and Cloud computing
Editor: Cristy Jane Burne <cristy.burne@gridtalk-project.eu>, Contributors: Marc-Elian Bégin <meb@sixsq.com>, Neil Geddes, Fotis Karagiannis Reviewer: Ignacio Martin Llorente <llorente@dacya.ucm.es>, Peter Kunszt <peter.kunszt@cscs.ch>, Malcolm Atkinson

Section 4, Security: A holistic approach
Editors: Fiorenzo Scaroni <fiorenzo.scaroni@switch.ch>, Christoph Graf <christoph.graf@switch.ch> Contributors: Christoph Witzig <christoph.witzig@switch.ch> Reviewers: Neil Geddes, Christos Kanellopoulos <skanct@grid.auth.gr>, Diego Lopez <diego.lopez@rediris.es>, Milan Sova <sova@ces-net.cz>, Daniel Kouril <kouril@ics.muni.cz>

Section 5, Service-centric e-Infrastructures through virtualisation
Editor: Radek Januszewski, Contributors: Mauro Campanella <Mauro.Campanella@garr.it>, Fotis Karagiannis Reviewers: Ludek Matyska <<ludek@ics.muni.cz>, Jesus Marco, Mauro Campanella

Section 6, Remote Instrumentation
Editor: Norbert Meyer <meyer@man.poznan.pl>, Contributors: Marcin Lawenda <lawenda@man.poz-nan.pl>, Reviewer: Jesus Marco <marco@ifca.unican.es>, Costas Cotsokalis <Constantinos.Kotsokalis@udo.edu>

Section 7, Sustainability of the computing-related e-Infrastructure
Editor: Fotis Karagiannis, Contributors: Mirco Mazzucato <mirco.mazzucato@pd.infn.it>, Peter Kunszt <peter.kunszt@cscs.ch>, Thomas Eickermann <th.eickermann@fz-juelich.de> Reviewers: Neil Geddes, Sergi Girona <sergi.girona@bsc.es>, Rosette Vandenbroucke <vandenbroucke@helios.iihe.ac.be>, Peter Kunszt, Gabrielle Von Voigt <vonvoigt@rrzn.uni-hannover.de>

Extensive background information, including known projects and experts, other actors, and some starting points is available upon request from the e-IRG secretariat.

8. *Annex I - White Paper Terms of Reference (WP ToR)*

Audience	Value provided	Abstraction level	Method for evaluation
Policy makers	Overview of technology-based opportunities (e.g. related to Lisbon objectives) and policy issues that need to be addressed on European level to get the benefits from them	High, policy issues need to be presented on "use case"-level	Provides ready, concise answers to a hypothetical question: "why should we pay attention?" from ministry level (e.g. one appointing a delegate) - or the press.
e-Infrastructure service providers	Help in estimating current and future requirements for "qualifying" to join in the common e-infrastructure (and thus ensure maximum return on investments in the service deployment and development)	High to medium, relative abstract check-lists for internal assessment	Is it possible for different persons to use the document and come up with similar assessment of the production readiness of a hypothetical service.
User communities	Overview of the possible services provided and the level of organisational structure and other resources the community needs to develop in order to get access	"use case" & check-list level	Can the White Paper be used to distil goals for a "Birds of Feather" meeting of a new user community

The format of the White Paper should follow the following guidelines:

- The main part of the document should be no longer than 10 pages (not including appendixes);
- The final list of topics will be confirmed to the e-IRGSP2 by the e-IRG chair, after the consultations required by the e-IRG rules and practices;
- The document should "make the case" for the policy recommendations on its own, without relaying on the appendixes.

The rationale for this format is to ensure that the best practices and consensus opinions are presented in a format that maximises the likelihood that they will be taken into account when making higher level policy decisions that form the regulatory framework within which the different actors in the e-Infrastructures operate in.

8.1. Communication methods

The White Paper editor appointed by the e-IRG chair will be leading the work and will report to the e-IRG board either through the e-IRG chair or directly, when the importance or urgency of the issue requires it. The e-IRGSP2 will complement this work by:

- Copying the e-IRG board in the messages related to the policy-level decision making in the White Paper and any contributions produced at the request of the White Paper editor;
- Summarising the developments regarding the White Paper in minutes of the biweekly excom telephone conferences.

Unless otherwise agreed for specific subtasks, the White Paper editor will copy the secretariat –mail address in the communication with the e-IRGSP2.

8.2. Expert list & initial reference material

The list of experts that can be consulted will be based on:

- List of expert produced by the e-IRGSP project[36], which includes Contributors to the previous versions of the White Paper and the Roadmap;
- Representatives of national and regional e-Infrastructure projects[37].

Other experts that the e-IRGSP2 wishes to consult need to be discussed and agreed with the White Paper editor (with a copy of the communications sent to the e-IRG board).

The background material that will be used in the preparation of the White Paper includes the earlier e-IRG documents, including workshop reports, workshop presentations, White Paper and roadmaps. The documents that have been sent to the e-IRG delegates or presented in the closed e-IRG meetings can be used if they are publicly available. Other background material will be discussed with the White Paper editor in a same manner as additional experts.

8.3. Background research strategy

The initial research strategy will be based on searching for updated versions of the background material that has been included initial reference material. Findings that contradict the earlier material – either in the new versions of the background material found or through other "due diligence" steps – will be immediately brought into the attention of the White Paper editor and the e-IRG board.

Additional material can be researched if agreed with the White Paper editor. These inclusions and their use will be reported in the excom telephone meetings.

8.4. Delegation strategy

Delegation of tasks by the White Paper editor will be documented in the excom telephone conferences.

36 Available at: http://www.e-irg.eu/support1/list-of-experts/

37 NRENs and initiatives mentioned in the EGI knowledge base (http://knowledge.eu-egi.org/index.php/Main_Page)

9. *Annex II - White Paper workplan White paper work plan 2008 - 2009*

V1.0, 21-11-2008, Matti Heikkurinen, Fotis Karagiannis, Marie-Christine Sawley

9.1. Introduction

The purpose of this document is to propose a general framework for the different phases of the e-IRG white papers: initiation, topics, workplan, timeplan. It also offers and a concrete implementation plan for 2008-2009. Currently, this implementation is planned for the whole year.

The work plan complies with the White Paper Terms of Reference (ToR) document that has been produced by e-IRGSP2 and discussed at the e-IRG Board level. This document intends to add some crucial details to the ToR that are needed for managing the White Paper project.

9.2. Input needed

The **main decisions required for** the e-IRG white paper (WP) process can be summarised as follows:

1. e-IRG white paper cycle (e.g. 12 months assumed by the e-IRGSP2 DoW);
2. Sources for white paper themes (usually e.g. the e-IRG workshops and task forces + approved background material; see Annex I for the workshop template);
3. ”Specifications” of the White Paper: number of pages, number of topics, aspects to be highlighted etc.

Each e-IRG WP plan is accepted when the Work plan and Time plan, as described in this paper, are approved by the e-IRG Board. The Board can decide to bring the decision forward to the e-IRG plenum.

White paper boundary conditions
A list of boundary conditions is proposed for any white paper. The boundary conditions should be in line with the Terms of Reference, which include mainly the purpose and format of the document, the communication methods, the expert list and the referenced material, as well as the background research strategy.

1. **Context:** The white paper should tackle policy-related topics and each section should be self-consistent, not relying on external sources or appendices;
2. **Size:** It is important that the white paper is concise, i.e. not more than 2 pages per section;
3. **Template:** A voluntary template can be used to gather input for the e-IRG workshops (in case of multiple parallel sessions), the white paper contributions by the section editors, the presentations / reports from the invited speakers or the task forces etc. An exemplary template can be found as Annex II;
4. **Annexes:** Technical or other related material should be kept as annexes at the end of the document;
5. **Terminology:** A terminology section should be also kept as an annex.

10 Available at: http://www.e-irg.eu/support1/list-of-experts/

11 NRENs and initiatives mentioned in the EGI knowledge base (http://knowledge.eu-egi.org/index.php/Main_Page)

9.3. Work plan

The **main steps in the e-IRG white paper process** can be summarised as follows:

1. **Choice of the WP topics and sources:** the main e-IRG stakeholders, such as the rotating presidency chair, permanent chair, e-IRG board, task forces, e-Infrastructure projects, e-IRGSP2 are the sources for White Paper topics. The e -IRG board and plenum decide on the priorities to give to the topics, in case not all of them can fit in the agreed WP cycle;
2. **Choice of the WP editing procedure** by the e-IRG board in cooperation with and the support of e-IRGSP2. Different procedures for filling the sections are possible:
   a. WP sections are based on a plenary workshop with invited speakers. The invited speakers of the workshop become the section editors and the main editor integrates the contributions in a coherent way;
   b. WP sections are produced on a full-blown workshop with parallel sessions for each topic and rapporteurs from each session. Session workshop participants can become contributors to each section, supported by other experts that e-IRGSP2 can identify with a specific process;
   c. WP sections are based on the outcomes of the task forces. The task force coordinator becomes the section editor and the task members the contributors (as experts).

   Other approaches can be identified, such as collaboration with e-Infrastructure projects, concertation groups and other external task forces. Obviously a combination of the above options is also possible.

3. **Appointment of a main editor by the e-IRG chair** (in consultation with the e-IRG board, if necessary) and defining the "Delegation strategy" (as specified by the ToR). For example:
   a. The e-IRG chair needs to be informed about appointment of section editors and (if required) experts for each topic;
   b. In addition, the e-IRGSP2 recommends appointing an independent reviewer, who will provide a review report to the e-IRG board for each major e-IRG White Paper draft versions;
4. **Editing iterations:** the content is based on the edited material from the sources (e.g. workshop and task forces reports). Iterations of the white paper edition are performed by the main editor and the section editors supported by e-IRGSP2 and reviewed by the e-IRG structures. If the second presidency within a year injects some new topics some steps might have to be repeated;
5. Iterations of the WP edition by the main editor and the section editors supported by e-IRGSP2 and reviewed by the e-IRG structures. If the second presidency within a year injects some new topics some steps might have to be repeated;
6. **Final version:** The pre-final version is proposed by the e-IRG board to the e-IRG plenum. e-IRG provides its feedback and the final editing iteration produces the final version;
7. **Recommendations:** e-IRGSP2 produces a list of recommendations which should be reviewed by the e-IRG board and pushed forward to the e-IRG plenum for approval. e-IRG feedback and final editing/approval iteration;
8. **WP publication:** The white paper and the recommendations are published, possibly as separate documents.

9.4. Time plan

1. **Decisions:** The main decisions which identify the directions taken (on WP cycle, sources for white paper themes and layout) should be agreed upon at the end of the previous presidency when the WP cycle begins. If this is not the case, it should be done in the first meeting under the new presidency. For 2008, this has been done during the first semester, through the consolidation of this document;
2. **Main steps:** The main steps should also commence in the first year semester. This holds for the topics, the workshop layout, and the designation of the section editors. A lapse of two months is proposed between the workshop and the last e-IRG meeting under the former presidency, while one month is given for the e-IRG feedback on the recommendations. Note that this plan is proposed for a single WP per year.

9.4.1. WP2009 schedule and milestones

1st semester

- e-IRG White paper source / layout: Workshop with invited speakers in a plenary mode, 24-25 April 2008;
- Invited speakers providing input on their topics: 26 May 2008;
- Completion of the list of topics (after consultations): 30 May 2008;
- e-IRG white paper editor: skeleton (topics + structure; text if possible) for e-IRG board review: 9 June 2008;
- e-IRG white paper editor: 2nd skeleton for e-IRG plenum comments: 13 June 2008 (2 weeks before Lugano meeting);
- Discussion in Lugano: 27 June 2008
    1. Decision on what are the topics that are certain to be in
    2. Strategy for proceeding with these topics (during the Summer);
- Comments by e-IRG plenum: 11 July 2008.

2nd semester

- e-IRG white paper editor: 1st draft of the White Paper (text on the topics agreed in Lugano + comments on the "candidate topics") for plenum: September 2008;
- Review of the list of topics, strategy for completing the document (incl. linkage with the Workshop topics): proposal from the e-IRG white paper editor to the e-IRG board: 26 of September 2008;
- Next version to be sent to the e-IRG plenum: 10th October (two weeks before Paris meeting);
- Feedback at the e-IRG meeting: 22 October 2008;
- Assignment of reviewers: 27 October 2008;
- e-IRG white paper editor: 2nd version for e-IRG board review: November 2008;
- Reviewers input: 10 November 2008;
- For e-IRG plenum review (November 24th - 2 weeks before Paris meeting) ;
- Discussion at the Paris meeting: 8th December 2008;
- Comments by e-IRG plenum: 2 weeks after Paris: 22nd December 2008;
- e-IRG white paper editor: 3rd version 4 weeks after Versailles: January 9th;
- Public consultation: 6 weeks (based on a live document, starting from the "Paris-version", to be updated in a transparent way);
- Recommendations: Start work in 2009.

Annex I – Workshop Template

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

Policy Area: A short paragraph description of the policy topic under discussion

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

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

Possible questions to be answered during the workshop:
Q1
Q2
…

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

Recommendation: This section should be prepared following the results of the workshop and the White Paper editing procedure. A list of distilled recommendations for each topic should be put forward, provided that there is a mature proposition. If a recommendation is available at an earlier stage, this should be put forward to the audience for comments.

Annex II – White paper contribution template

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

Policy Area: A short paragraph description of the policy topic under discussion

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

Context: Current practices - achievements and limitations: A short section describing recent advances and limitations in the area i.e. policy state of the art; more technical material should be placed in an appendix (if needed). It should highlight the main issues with the current situation that need to be tackled.

If during the workshop presentation a list of challenges were presented in the form of questions, and the questions were answered, these can be summarised in the workshop report

Possible answers to the questions that rose during the workshop:
Q1/A1 (Question 1/Answer 1)
Q2/A2
…

Proposed approach: In case there is a foreseen solution or proposal under consideration, this should be presented in this section.

Recommendation: This section should be prepared following the results of the workshop and the White Paper editing procedure. A list of distilled recommendations for each topic should be put forward, provided that there is a mature proposition. If a recommendation is available at an earlier stage, this should be put forward for comments.