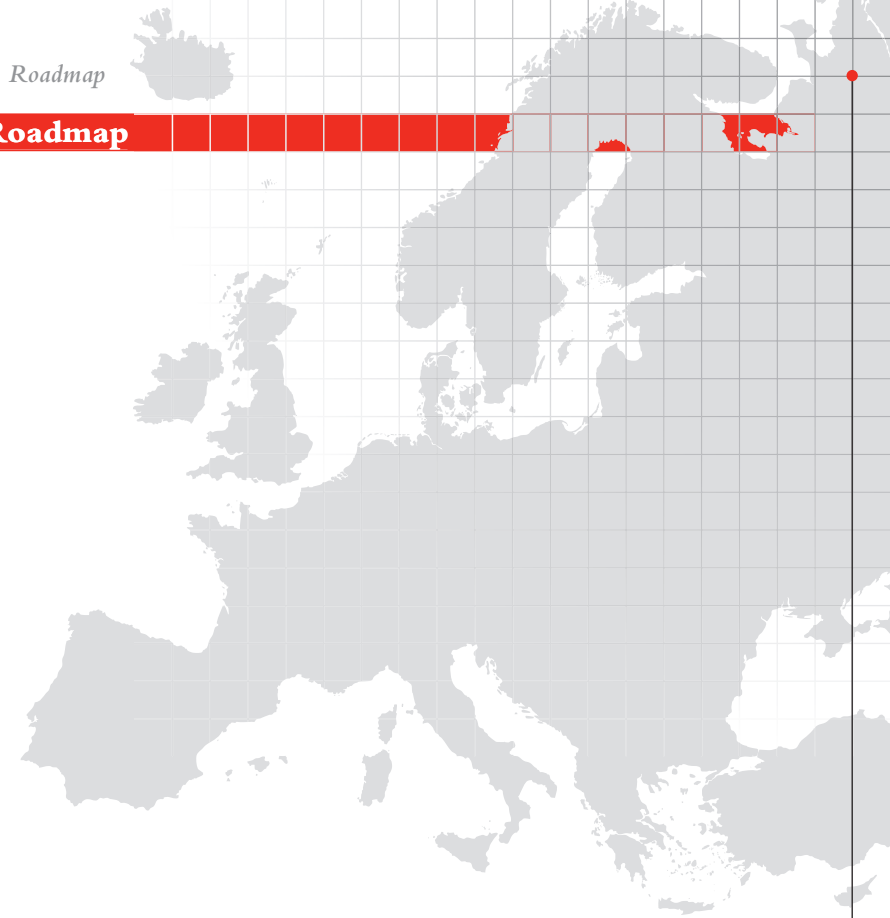




e-IRG
e-Infrastructure
Reflection Group

<http://www.e-irg.eu>

e-Infrastructures Roadmap 2007



e-Infrastructures Roadmap

*Leenaars, M.A.G.J., Karayannis, F. et al (ed.)
Second edition, October 2007.*

*Published by e-IRGSP
P.O. Box 93575
2509 AN The Hague
The Netherlands*

*Tel: +31 (0)70 3440 526
Fax: +31 (0)70 3440 946
Email: secretariat@e-irg.eu
Web: <http://www.e-irg.eu/support>*

Designed by No Panic communication services, www.nopanic.nl.

Printed in The Netherlands.



This work is licensed under a Creative Commons Attribution 4.0 International License.



Table of contents

Preface	6
Introduction	8
What is an e-Infrastructure?	9
Why should Europe join forces?	12
How to read this Roadmap	16
Networking Infrastructure	18
Global End-to-End Hybrid Networking	19
Middleware infrastructure and organisation	21
Authentication and authorisation infrastructure	22
Software life cycle management	24
Middleware repositories and parameter registration	27
Ensure open standards	29
Training & support for scientists and support personnel	30
Resources	32
Supercomputer infrastructures for Europe	33
European Storage Facilities	35
Service Oriented instruments and facilities	37
Sensor networks	39
Computational accelerators	41
Incentives for providing grid resources	43
Leveraging new technologies	44
Data handling	46
The Knowledge Life Cycle Infrastructure	46
Crossing the boundaries of science	49
Scientific collaboration	50
Working together with industry	51

Preface

This document is an updated version of the e-Infrastructures Roadmap produced by the e-Infrastructure Reflection Group in 2005. After the European Commission proposed in 2004 to have a strategic roadmap for Research Infrastructures developed to cover the next 10 to 20 years, e-IRG accepted the challenge to draft a roadmap for the e-Infrastructure area. e-IRG focuses on the genuine general IT infrastructures that need to exist irrespective of disciplines or specific sciences - encompassing networking, highly advanced computing and grids and storage. e-Infrastructures are by nature supportive service activities, they are the tools our scientists and researchers work with. The first e-IRG Roadmap was approved by the e-IRG delegates in December 2005 in London.

In order for the roadmap to be a reliable policy instrument, it was found necessary to regularly update the Roadmap, to reflect new trends, technologies, insights and initiatives. Providing an adequate e-infrastructure means operating in a field with very rapid changes. So when the roadmap was made publicly available, the work on the first update was started almost immediately. A consultation under the stakeholders identified in the Roadmap was held, output was gathered during a number of e-IRG workshops and input was sought from representatives within the European scientific community. This work would not have been possible without the broad support and contributions from a large number of people. We are grateful for your contribution but due to limited space we cannot mention by name everybody involved. Thank you!

The result is the revised roadmap you have in front of you. Important changes include the elevation of data infrastructure to a higher level (Data Handling) and a separate section on „The Knowledge Life Cycle“, increasing the level of detail in this area. There is also a new section on computational accelerations. And many changes were made at a more detailed level, taking into account the start of FP7, developments within ESFRI, new technologies and more.

The 2007 e-Infrastructures Roadmap gives the reader a long term view of the development and the foreseen trends in the usage of e-Infrastructures to support the European development and welfare. Europe needs to collaborate to be effective and offer a European competitive advantage. This means we need services for improved access and sharing of research information resources across research discipline borders (primary research outputs and data). New generations of researchers will have a different perspective on, and requirements of, the research process due to increased familiarity and use of ICT tools. New paradigms and fresh ways of thinking should be facilitated at the e-infrastructure level. In parallel we need many new skills to be able to utilize fully this new landscape, which need to be developed in parallel to the more physical infrastructures. These and many more challenges are outlined in the 2007 e-Infrastructures Roadmap. We hope you will find it both inspiring and useful, and look forward to your response.

Dr. Leif Laaksonen
e-IRG Chair

Drs. Michiel Leenaars
Lead editor Roadmap

Introduction

In the 20th century Europeans produced the first computer¹, invented packet switching² (the basis for the technology operating the internet) and more recently conceived the world wide web³. However, Europe has fallen behind in reaping the benefits of its own innovative force. Now, with grid technology as a strong catalyst and with the parallel deployment of a world leading networking infrastructure Europe has a major chance to regain its former leading position. The e-Infrastructure is seen as the spinal cord of the European Research Area delivering advanced facilities driving the testing and the first deployment of new innovative technologies.

The current technological leadership must be capitalized on while there is still very much a green-field situation that allows us to grasp the new research and business opportunities. Europe must accept the challenge to develop and build the e-infrastructure required for the information age now – an investment opportunity not to be missed. With new contenders – like the fast-growing Asian economies – already looming on the horizon, Europe needs to seek the front ranks again if it wishes not to be marginalized in due course.

In the next decades science and research will change fundamentally in the way they operate^{4,5}, so the scope of thought should surpass the current situation and needs. In order to support Europe-wide communities that are able to interact in a global environment as equals, it is important to encourage sharing of electronic infrastructure resources as a way to create suitable conditions for cross-disciplinary interaction, providing fertile ground for innovation and eventual industrial exploitation and use in education.

There is no doubt also that the impact of new infrastructures will be far beyond science, as was witnessed with both the internet and the world wide web. Possible uses of the new infrastructures outside of the research and education communities include commercial services, security and disaster management, digital libraries, entertainment (digital television, rich media distribution, gaming) and e-learning. Enhanced competitiveness in these areas positively impacts vast parts of the European economy and offers tremendous opportunities. Collaboration and information exchange with industry – both as supplier and as a user community – and the rest of the globe is necessarily a part of the entire approach. Of course combining the major efforts from the research area and those from industry will be of great help to create a mature and sustainable market through orchestration and convergence of competing and complementary technologies⁶.

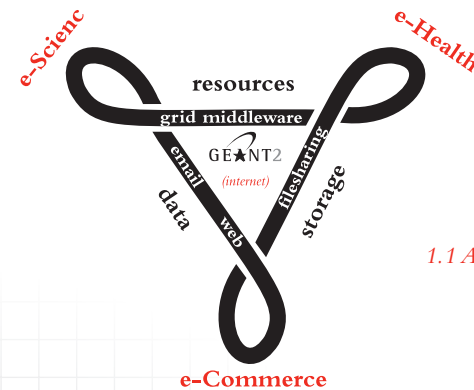
1 Konrad Zuse, 1936. See: Rojas Raul (ed): *Die Rechenmaschinen von Konrad Zuse*, Springer Verlag, 1998.

2 Donald Watts Davies, 1965. See: D.W. Davies, K.A. Bartlett, R.A. Scantlebury, and P.T. Wilkinson, "A digital communications network for computers giving rapid response at remote terminals," *ACM Symp. Operating Systems Problems*, Oct. 1967.

What is an e-Infrastructure?

The term e-Infrastructure is used to indicate the integrated Information and Communication Technologies (ICT)-based Research Infrastructure in Europe. Of course such an infrastructure builds on many ICT components that have been around for quite a while, such as networks, supercomputers, and storage. There are many interdependencies between these components, so their future should be planned coherently. The e-Infrastructure viewpoint allows to join and fit all interrelated infrastructures together and start think of them as a system – and optimise not for each individual part, but for the whole. The prime goal of the e-Infrastructure may be to support e-science, e-health and e-culture, but at the same time opportunities are created for many other application domains that contribute to society such as e-commerce, e-government, e-training and e-education.

A competitive e-infrastructure is indispensable for the oriented branches of the sciences. Well-known examples are climate and earth system research, water management, fluid dynamics, biophysics, theoretical chemistry, astrophysics, quantum chromodynamics, nanostructure physics⁷ and high-energy physics. Both the increasing progress on mathematical models and the complexity of simulations cause the demand of these subject areas for computing cycles to be almost without limits. But also from traditionally less computer-oriented areas such as the social sciences, the humanities and biodiversity there is a strong trend towards mass deployment of ICT to manage the large variety of decentralised data sources and find novel approaches to traditional problems.



1.1 A schematic overview of e-Infrastructure components.

3 Tim Berners-Lee, 1989. See: <http://www.w3.org/People/Berners-Lee>

4 Ian Foster, 2005. *Service-Oriented Science*. *Science*, vol. 308(5723), 814 – 817.

5 Tony Hey and Anne E. Trefethen, 2005. *Cyberinfrastructure for e-Science*. *Science*, vol. 308(5723), 817 – 821.

6 However, it is not a goal in itself. Industry is made up from a large amount of autonomous actors that cannot possibly all be involved in the same level with every development, and much happens in parallel only to be decided by the market. So while a broad coalition of industrial partners can be a decisive help in the successful adoption of certain technologies, the future remains as unpredictable as ever. Even if market leading industrial partners are heavily involved throughout the entire process in a fundamental way, other technological solutions from smaller or even unknown competitors might well prevail – even if they are technologically inferior and/or incomplete.

7 For more examples see: *Recommendation on the Installation of European Supercomputers*, Wissenschaftsrat, november 04

Key components of the e-infrastructure are *networking infrastructures, middleware and organisation*, various types of *resources* (such as supercomputers, sensors, storage facilities) and data. In diagram 1.1, the relationships between those components become clear. The network is at the heart of everything, and in the age of hybrid networks includes lambda-networking (popularly known as ‘light paths’). The middleware and virtual organisations⁸ (exemplified by the black knotted triangle) connect the distributed resources, storage facilities and data in a seamless way. The data infrastructure is not focused on real-time availability but on deriving, adding and correcting distributed knowledge. The application domains (such as e-science and e-health) are on the outside of the chart to exemplify the parties served by the infrastructure; these are only relevant insofar as they bring in resources.

The four components are:

- **Networking Infrastructure**

The research networking infrastructure delivers the physical connections for the e-Infrastructure. These will primarily be delivered through the hybrid GÉANT pan-European backbone network and the fine-grained National Educational and Research Networks. These networks together form the solid basis for the general purpose scientific communication, supporting collaboration and special uses (of which the grid and distributed supercomputing applications are but a few). Europe should continue to play the leading role in building both European-wide and international global end-to-end connectivity for the research and academic community. Additionally, when linking to resources that are outside the scientific domain (such as public utility, commercial or military resources) other networks may be added.

- **Middleware and organisation**

Middleware plays the intermediary role to facilitate a deep integration of individual components with the networks into a European Science Grid. Grids are an evolutionary step in the way we can work with computers and everything connected to them. A grid consists in principle of a group of resources (digital devices and anything attached to them or stored on them) which can be used for combined efforts. The middleware assumes a network such as the internet to run on – or in the case of GÉANT a hybrid IP/optical network – and in fact implements a protocol stack into an interoperable runtime environment and/or query mechanism that allows for sharing of information and tasks between distributed devices and systems. New processes and procedures have to be devised to alter the way organisations work, delivering for instance an authentication and authorisation framework. In order to enable a separation from generic facilities and discipline-oriented solutions special attention is needed to support and train users.

- **Resources (such as supercomputers, sensors, data)**

The European Science Grid as an integrated approach to serve the European scientific user communities should be populated with a number of resources in order for it to add value to the individual components. The word resources in this context should be interpreted in a broad way, covering literally everything that is of interest to science from computers, large storage facilities, telescopes, satellites, special physics equipment, weather balloons, lasers, spectrometers, visualisation means and large sensor networks. A resource can also refer to artificial intelligence agents and even people as support organizations that can be shared between institutes. The only requirement is that the resource (from supercomputer to cellphone) can at some point exchange the necessary information through standardized interfaces, i.e., grid protocols. The end goal is a rich ecosystem of resources that offer a broad gamma of hardware, software, services and data spaces.

- **Data**

A special type of resource is data. Data fuels the information age, and databases and digital libraries are the binary oil fields of that age. The amount of data is growing at an extremely high pace in what has become commonly known as the data explosion⁹ or data deluge.¹⁰ Handling the continuous shock waves of new scientific data output and serving bursty usage patterns efficiently, while keeping that data reliable and available at the fingertips of the scientist and the historian, can only succeed with a proper data life cycle management infrastructure that leaves the distributed nature of the scientific process intact.

The e-infrastructure cannot be seen independently of other more generic advanced ICT collaboration tools – such as sharing remote work spaces, advanced presence services and high resolution video conferencing and visualisation. These will develop in their own right, but they are essential when aiming for low-level integration of communities around e-infrastructures.

⁸ A virtual organisation as meant here is a partnership of organisations electronically sharing resources, often for a common goal or with a common interest.

⁹ Hillary Jay Kelley ‘Philosophy of Science’, Vol. 36, No. 2 (Jun., 1969), pp. 178-196.

¹⁰ Hey, A. J. G. and Trefethen, A. E. (2003) ‘The Data Deluge: An e-Science Perspective’, in Berman, F., Fox, G. C. and Hey, A. J. G., Eds. *Grid Computing - Making the Global Infrastructure a Reality*, chapter 36, pages pp. 809-824. Wiley and Sons.

Why should Europe join forces?

It takes a lot of effort to try to combine resources scattered across the continent. Why do we take all this trouble? The answer is simple: there is a long term huge structural need for a variety of resources throughout all scientific domains that can be best satisfied that way. It is a given that we already have many resources deployed. With an e-Infrastructure in place we can use them smarter and far more cost efficient at the same time – thereby giving a higher return on investment and increasing their potential. Adopting the grid paradigm will allow us to think global and act local.

Below we will try to give a number of reasons why an integrated approach to ICT infrastructure is preferable over others. One such reason is that investments on the national scale have proven to be insufficient to provide world class resources to European scientific communities even though total European expenditure is up to par. Another is that the sustained demand for resources on a European level requires a structural approach that solves the huge inefficiency that pressures individual science domains to justify their own facilities over and over again. Another is that a European shared approach will yield a faster *time-to-science*¹¹ because of pooling resources while lowering prices through just-in-time acquisition.

National scale is too small

Even though the European economy as a whole is very strong indeed, individual European countries cannot sustainably provide world class resources to their scientific communities by themselves. An area where this is clearly visible as a trend with data available over a long time is supercomputing. There has been a significant gap between the fragmented European supercomputer facilities and what leading sites in Japan and the USA have had available for the last decades. Computers at the very top of the market cannot be funded by a single European country; they have to be supported by multiple nations and need to be co-financed on a European level. There is an urgent need for the European Commission and the member states to review and bridge that gap. Similarly, there is a lack of variety of supercomputers from the point of view of architectures that are available in Europe. Scale gets even more important if Europe is to undertake some grand challenges: large scale scientific endeavours that can push the boundaries of science significantly further. A European focus can make such undertakings happen.

Sustained growth in volume requires a structural solution

When one looks at the figures, the overall picture is clear. The investments in e-Infrastructure components in Europe have relentlessly risen in the last two decades – without a single dip. The systems and networks may have gotten much cheaper, but demand has been growing even faster. Some key figures: investments in networking are now at about 0.4 billion Euro annually for GÉANT and the networks run by NRENs alone (so excluding the fine-grained ‘campus-level’ deployment which is the most expensive). The current installed base for supercomputers and very large clusters in Europe (200 largest European computer systems currently in place) is worth over 0.6 billion Euro¹². This is without operational costs such as maintenance or power consumption taken into account, and at current price levels – of course what was actually paid for these systems at the time of purchase will have been a multiple of that amount. In the EU FP6 0.65 billion Euro was spent on large Research Infrastructures,¹³ of which significant amounts go to ICT components. A single sensor network alone can have a budget of 0.15 billion Euro, of which about one third may go to special supercomputing facilities.¹⁴

Clearly, the sustained nature of these investments and the technical overlap among many of these activities justifies coordination and consolidation. It is very inefficient if every scientific domain that requires some large facility has to find justification and political leverage to obtain funds, and then built the expertise to make the right choices for huge one time investments.¹⁵ Such a procedure costs a lot of time and money, which also makes that emerging sciences are at a disadvantage. Once investments are made, new insights may dictate starting all over again. Centralised facilities that are instantly available at a fair price without the need for everyone to go through the whole operation will make science more effective.¹⁶

¹¹ ‘Time-to-science’ indicates (equivalent to the commercial wording ‘time-to-market’) the time between the start of a scientific research endeavour and its final publication. Used in: Space Studies Board (SSB) Engineering and Physical Sciences (DEPS), “Assessment of Mission Size Trade-offs for NASA’s Earth and Space Science Missions”, NASA 2000.

¹² Estimate derived from the e-IRG Knowledge Base, <http://knowledgebase.e-IRG.eu>

¹³ See: <http://www.cordis.lu>

¹⁴ See: <http://www.lofar.org>

¹⁵ An example of collaboration on a European level is the European Centre for Medium-Range Weather Forecasts (ECMWF), which is a co-operation of 24 European national meteorological institutes. ECMWF operates a state-of-the-art high-performance computer, which always belongs to the top systems in Europe.

¹⁶ In the case of supercomputing facilities commercial offerings from companies like IBM, SUN and Amazon in this area have recently seen the light. If the case for a specific type of resource cannot be made, one could look at the market to supply it.

Lowering time-to-science and just-in-time acquisition

With investments in ICT timing is a key issue: it is clearly often disadvantageous to be too early as investments in ICT hardware have to be written off very fast. In the case of supercomputers: in the last two decades it has taken about seven years for the fastest computer in the world at a certain point in time to disappear from the top 500 of available systems – in fact being surpassed by cheap new systems that cost less than 1% of the original price. Every month that such a system is not fully utilised during its first years of deployment is extremely costly indeed – with a price tag in the order of magnitude of one to five million Euro per month.¹⁷ Any technological glitches are often the risk of the buyer. Deployment of immature technology is not only very costly in this respect, but can be very time-consuming for the scientists involved – with a slowdown of development and decreased enthusiasm and support from the scientific community as a result.

However, it is equally dangerous and demotivating to be too late: if scientists from inside the EU have a structural competitive disadvantage through lack of appropriate resources their research will suffer and they will probably miss out on the commercial spin-off entirely. Global R&D expenditure on a computationally intensive research area as pharmaceuticals alone is estimated to be over 40 billion Euro per year, with an average increase of 5.4%. The development of new medicines takes an average of 12 years.¹⁸ A medicine can only be patented once, so the stakes are high in the race to be the first one to apply for the patent – the outcome of which will make the difference between a huge profit or a huge loss.

Time-to-science (from proposal to deployment) for a very large resource takes at least half a year – and often much longer. So waiting until the science is ready would in any case waste valuable years. Making sure the needs of the researchers and the timing of investments are optimally aligned is a strategic activity with significant risks of ending up with very expensive capacity too early while lacking vital capacity later on. The risk to be just in time to invest in exactly the right technology can easily be reduced once an e-Infrastructure is in place.¹⁹ If there is a pool of resources that is shared at a European level, allocation of resources can take place much faster – significantly lowering time-to-science.²⁰ If there is a need for more resources than what is available, these can be constructed just-in-time. Because of the larger overall volume of facilities individual projects with urgent requirements can be facilitated faster.

Enabling a European Science Grid

In this Roadmap we use the term European Science Grid (in singular, even though it would arguably be more accurate to use the plural) in order to describe the pool of resources that are somehow available through any of the grid environments running at a given moment in time and that are predominantly driven by European incentive – along with the technological and organisational efforts that are in place to bring them together. The vision of a single monolithic integrated European grid environment is probably no more desirable than a similar monolithic integrated global grid environment, as it will lead to a bloated and hardly usable facility. We can imagine the outlines of such a concept if we think of what enormous opportunities would arise if we could create a grid of grids, or a virtual meta-grid able to interconnect all the grids in Europe (and probably before that having already integrated many across the whole planet) together without actual low-level integration.

Why choose the boundaries of Europe as the appropriate level to optimize? There are of course other criteria one can sensibly use to single out a group of grids, such as disciplinary boundaries or national or smaller regional boundaries at which major funding is shared. Such boundaries are meaningful and have their own economic reward in increased effectiveness – and in the case of national grids increased competitiveness among the member states. As such they will not need support from a European perspective to be sustainable – except perhaps in cases where help is required to help grid-nascent countries participate in pan-European initiatives. The European level is however in most cases more relevant since one needs critical mass and a strong mutual commitment. The European Union is a logical choice as it interconnects heavily on an economic level and has increasingly shared funding.

Grid technology is a major enabler for world-wide partnership and teamwork, catalysing interaction between major and minor players in the e-Science field and beyond, while providing new ideas and possibilities to apply best practices learned from others to develop services in each research centre. The European e-infrastructure will – through the European Science Grid and other tools it provides – support more intense collaboration between various research centres and their researchers than ever before. The European Science Grid will embed the concept of Europe strongly within its infrastructure as it will quite naturally strive towards mutual interdependency by promoting an architecture that optimises availability and cost efficiency of scientific facilities not on a national level but within the European borders.

¹⁹ With new technologies such as hybrid networking Europe is indeed leading the world, which makes Europe more attractive to invest in. An additional advantage of getting in at the right time is that one can profit from the partnership with industry that is always keen to prove their technology is ready for large scale deployment. The only way to do this in a sensible way is on a European scale.

²⁰ One such collaboration that pools resources is DEISA (Distributed European Infrastructure for Supercomputing Applications), which interconnects the individual systems from a number of supercomputing centers to form a distributed terascale supercomputing facility by using grid technologies.

How to read this Roadmap

Networking, middleware, resources and data handling as described above together make up the e-infrastructure. In the roadmap we will define for all these components from the view of Europe as one interconnected system what we envision should be done. In this respect we have identified a number of opportunities for Europe that together will lead to a world leading, affordable and cost-efficient e-infrastructure. There are also some boundary conditions that have to be met, such as partnering with industry to avoid loss of critical mass, increased control over our own IT systems (open source) and the development of collaboration tools that will allow true European collaboration.

The opportunities that we will present in the Roadmap are:

Networking infrastructures

- Global end-to-end hybrid networking

Middleware and organisation

- Authentication and authorisation infrastructure
- Software life cycle management
- Middleware repositories and parameter registration
- Ensure open standards
- Training & support for scientists and support personnel

Resources

Building a healthy resource ecosystem

- Supercomputer Infrastructures for Europe
- European Storage Facilities
- Service Oriented instruments and facilities
- Sensor networks
- Computational accelerators

Resource Attraction and Evolution:

- Incentives for Providing Grid Resources
- Leveraging New Technologies

Data handling

- The Knowledge Life Cycle Infrastructure

Crossing the boundaries of science

- Collaboration tools and environments
- Working together with industry

The developments in many of these areas are still going at an extreme pace in with increasingly higher stakes. It will remain challenging to make even parts of this agenda come through. However, it is necessary to keep competitive and avoid losing attractiveness to academics and businesses; Europe cannot afford to lose important businesses or create a brain-drain – which will continue to happen if people cannot perform first class science due to lack of infrastructure.

A journey that will take us two decades into the future cannot be fully planned, but we hope this Roadmap will inspire you and act as a strategic guide for the long term development of the European Community. Since this is a Roadmap, we have tried to provide you with as much guidance as we can. We will not only explain what the opportunities are, but also give clear directions on how to tackle every opportunity. Every opportunity follows the same format.

- ➔ We provide a description of why a certain goal is strategic and what the short term actions are that should be taken as soon as possible (what to do at the *'Next turn'* of the road).
- ➔ We then describe the *End destination*: where should we be in twenty years (or sooner, if possible).
- ➔ We then identify per opportunity a number of *Relevant policies, organisations, activities*. Who should at least be involved? This inventory is by no means final and probably the most prone to omissions. Please bear in mind these are not meant to try to exclude any organisation or initiative that isn't mentioned. Anyone who can contribute is invited to step forward and help Europe get the e-infrastructure it needs to take on the competition for the 21st century.

Networking Infrastructure

The next generation optical pan-European network platform GÉANT2 was launched in 2005. It integrates advanced IP-based routed services with lower layer manageable end-to-end optical connections for the support of e-Science initiatives (e.g. Grids, collaborative research etc.) These parallel networking flows will better serve the diverse requirements of the European Research and Academic community. Any future incarnation of the trans-European network GÉANT should proceed further in dynamic provisioning of production quality seamless connectivity – unless the need for dynamic bandwidth falls short, in which case dynamic bandwidth should give its position to fully meshed-type permanent bandwidth. The National Research and Educational Networks will have to provide the fine-grained connections to institutions and resources. Europe should continue to play the leading role in building both European-wide and international global end-to-end connectivity for the Research and Academic community. Additionally, when linking to resources that are outside the scientific domain (such as public utility, commercial or military resources) other networks may be added.

There is only one destination available for networking infrastructure:

- Global end-to-end hybrid networking

Global End-to-End Hybrid Networking

Global end-to-end connectivity is a key issue. In order for end-to-end services to work in the hierarchical European backbone consisting of the campus, national, regional and European spans, the European networking infrastructure should universally deploy interoperable protocols. In addition, a human network consisting of the corresponding network administrators should be formalised to exchange views and ideas. Europe has invested heavily in the next-generation IPv6 protocol currently deployed in the pan-European backbone GÉANT and the majority of national networks. Europe should keep this leading position and cooperate with other interested regions (e.g. China) to build upon its expertise.

The pan-European networking infrastructure should aggressively cover all the European member states with world-class connections before the end of the decade, and be prepared for new member states entering the European Union. If the necessary amounts of so called Dark Fibre are not available in time, action must be taken. This dark fibre vision includes areas of South-Eastern European countries (SEEREN) along with Belarus, the Ukraine and Moldova, in an effort to ease the digital divide in Europe.

For the sake of cost-effectiveness, building dedicated network of links between major European research resources would be a good investment. This is due to the very advantageous cost/performance ratio of dark fibre solutions, when compared to the current standard (routed wide-area network infrastructure).

Although outside the direct scope of the Research networks, further expansion of broadband and optical networking, ultimately covering the last mile to the doorstep of households, companies, governments, institutes and organisations is a key factor of the development of the e-Infrastructure in the long run. Research networks should be seen as a technology enabler and catalyst for the proliferation of ICT usage in Europe. Solving the many (often legal) issues on a European level concerning the last mile is crucial. But it is not just the physical transport layer that needs attention: if we wish to secure pan-European end-to-end functionality some form of national coordination of network solutions, Quality of Service and security on the campus level and to institutions need to be considered.

There are significant opportunities for mobile networking. Some steps have been taken to found a European organisation that further investigates these opportunities and seeks to deploy a pan-European mobile networking environment.

The use of dark fibre acquired from the “new market” implements a new model of “ownership” of the networking resource, as it decouples the provision of the network from bandwidth provision – and the related pricing – by traditional carriers. This opens a completely new and innovative perspective for applications (like grids), as the cost of bandwidth is no longer a serious bottleneck for network provision. Longer term strategic issues not directly dependent on current practices and cutting edge technologies must drive e-Infrastructure planning, including research & education networking. The emerging business model should resolve fundamental questions like ownership of infrastructures, sharing policies, foresight of capital investment, consequences of technology driven choices etc.

Next turn

- ➔ Define and develop standard protocols and interfaces for network control and management planes coping with a multi-administrative, multi-technology and multi-equipment domain environment
- ➔ Work on interoperability of the grid middleware with the above network control and management planes
- ➔ Solving the many (often legal) issues on a European level concerning the last mile is crucial

End destination

- ➔ A reliable high speed hybrid network covering all of Europe and providing global end-to-end connectivity

Relevant policies, organisations, activities:

- ➔ NREN's, TERENA, DANTE, NRENPC, IETF, ITU, IEEE, DG Information Society and Media, GLIF, GLORIAD

Middleware infrastructure and organisation

Middleware plays the intermediary role to facilitate a deep integration of individual components with the networks into a European Science Grid. Grids are an evolutionary step in the way we can work with computers and everything connected to them. A grid consists in principle of a group of resources (digital devices and anything attached to them or stored on them) which can be used for combined efforts. The middleware assumes a network such as the internet to run on – or in the case of GÉANT a hybrid IP/optical network – and in fact implements a protocol stack into an interoperable runtime environment and/or query mechanism that allows for sharing of information and tasks between distributed devices and systems. New processes and procedures have to be devised to alter the way organisations work, delivering for instance an authentication and authorisation framework. In order to enable a separation from generic facilities and discipline-oriented solutions, special attention is needed to support and train users. Security, privacy and integrity of information are an essential part of all these aspects, because without those the new infrastructure will not take off.

A number of destinations are available for middleware infrastructure and organisation:

- Authentication and authorisation infrastructure
- Software life cycle management
- Middleware repositories and parameter registration
- Ensure open standards
- Training & support for scientists and support personnel

Directions to: Authentication and authorisation infrastructure

In order to build the European Science Grid, resources from many different organisations, both public and private, will need to be combined into a coherent system. This inter-organisational sharing requires well-established, trustworthy, and judicially sound ways of authenticating and authorising access to all services that comprise this infrastructure. Europe has taken a leading role in building authentication and authorisation based on a federative framework. Such federations recognise local autonomy of each resource and allow the European infrastructure to leverage organisational, national and pan-European trust mechanisms. They consistently retain local control, which allows the infrastructure to remain lightweight – a key feature of such an infrastructure if it is to remain manageable. It also leaves enough degrees of freedom on the national level to accommodate different policies and legislative conditions.

Initiatives like the IGTF and in Europe EUGridPMA have made important contributions towards global identity trust interoperability. Yet, Europe will need to maintain and refine this strategy in order to create an open trust hub that allows the influx of new and wider communities: first among those will be large amounts of scientists but also Small and Medium-sized Enterprises (SMEs) and other interested potential contributors to the grid developments. Europe should establish a framework for collaboration that would lead to convenient, interoperable mechanisms for both authentication and authorisation that enables a single integrated view of all resources within their domain of operation for the user. Where relevant, privacy and confidentiality of identity, access control, and information content should be adequately protected against malicious or accidental exposure. Especially in the realm of authorisation and access control a wide diversity of mechanisms has been deployed. Some of these bind more tightly into organisations (including physical controls like access to buildings) or link to organisational databases, whilst other mechanisms allow for individuals to participate in communities irrespective of their home organisation (the model adopted by the virtual organisations in the Grid).

The resulting European framework should allow for several models to be made compatible and incorporate them in a trust hub, with an overarching federation to ensure consistency and classification of the information provided by the federation's participants on at least a European but preferably a global scale. This will allow for an open process of flexible consortium building in both the science and enterprise domains. The effectiveness and reliability of the authentication and authorisation infrastructure will be a key factor in the success of many future activities and may contribute significantly to our competitiveness on a global scale in the long term.

Next turn

- ➔ Build on and extend federation based authentication and authorisation infrastructures to support involving a growing and broadening community
- ➔ Where possible, identity provisioning should leverage national digital identity initiatives
- ➔ Support the establishment of frameworks able to integrate all the (nation- or community-based) AA federations, in the spirit of the achievements of the federation for authentication

End destination

- ➔ A scalable, reliable and cost-efficient authentication and authorisation infrastructure accepted on a global scale

Relevant policies, organisations, activities

- ➔ International Grid Trust Federation, EUGridPMA, GÉANT, TERENA/TACAR, Eduroam, EuroPKI, DG Information Society and Media, DG Health and Consumer Protection, DG Research, DG Legal Service, DG Internal Audit Service, DG European Anti-Fraud Office, ENISA, Liberty Alliance, CACert

Directions to: Software life cycle management

The e-Infrastructure would be incomplete without software. Scientific software acts as a skeletal framework for many scientific developments, as it implements incremental knowledge, approaches, algorithms, and models. Often, specific scientific codes are used by dozens or even hundreds of groups and thousands of scientists across Europe and the rest of the world – with a life span of sometimes several decades. In some areas where complexity has risen to the point that researchers cannot but depend on the validity and accuracy of the code the software is becoming the most important carrier of scientific insights. There are two realms of codes: the first is owned by privately or publicly held companies, and requires significant fees to use them. In a way these act as scientific publishers, incorporating the knowledge discovered by others into the software. This creates a number of problems, one of which is that one party decides where innovation can take place, which architectures are supported and optimised for. In a grid environment such features are undesirable, as they shrink the available set of resources. People also tend to have the same problematic relationship as with more classical scientific publishers: since many scientists are dependent on the codes they run for carrying out their research, the financial drain may lead the community to an exhaustion point. Current Intellectual Property Right solutions are not in the interest of science. The only exit scenario is to recreate the entire functionality from scratch as a community effort, which is as time-consuming as it is error-prone with codes that have been worked on for many person years and of which the internals of the software are unclear to outsiders and need to be reverse engineered. Even if some user communities of non-public codes would embark on such a scenario, this would still require more support for them for the next few years. For others, leveraging the buying power of the joint European users is essential.

The second realm of code is the software that is in the purely scientific domain. Many software is spawned rather organically from small-scale initiatives -- quite often individual PhD-projects. Many of these codes have however outgrown their initial use and development scope, and the creators fall victim to their own success. Individual research groups that are now “responsible” for maintaining the codes -- as good or as bad as possible -- take considerable pride in providing such a service to the community. But they cannot be expected to pay for the whole development for years or even decades by themselves. At some point the software reaches a critical stage for their user communities: how does it continue to develop? Implementing other people’s algorithms, fine-tuning and optimising the code for new user groups and software environments, and debugging – in short keeping the code up to date and usable to all – lack the imaginative force of new discovery. Despite the importance of such codes and their widespread use, they often lack adequate development support or even a basic life cycle management infrastructure. It is not that people are not seeing the importance of maintenance,

optimisation and further development. On the contrary, but such activities are outside the scope of the basic funding structure and national orientation of most research funding agencies. In fact, software costs – including those of commercial codes – are often systematically off-radar. That it is problematic is very clear, when we are about to enter an era where e-science is to blossom. Solutions at the EU level are needed; software crosses boundaries and local solutions are inefficient.

There are three scenarios: the first is to leave the software as is, and just let it loose on a large scale on the e-Infrastructure. This will at best just maintain inefficiency, and thereby invisibly occupy resources equivalent to significant amounts of money. If no investments are made, the science behind the software will suffer. At worst the community suffers from an accumulation of systematic programming errors, resulting in large scale scientific errors undetectable until someone does build another correct implementation or version of the software. This is obviously penny-wise and pound-foolish.

The second scenario is to commercialize all the individual software packages. This may pay for some hard-needed initial changes and be economically sustainable in the future, but it is both hard to do from a legal point of view (since so far everything was built with public funds and with the contributions of many) and would create other problems. Commercialising may for instance take away the code from the open source domain – which is important for progress as it allows scientific scrutiny and the discovery of errors on the one hand and enables innovative dispute on the other. And of course it will bring about the problems described earlier as the software enters the commercial realm: reducing the amount of architectures and software environments supported and delivering scientists to the mercy of their software publisher.

A third option would be to create some structural funds in order to encourage professional software engineers and scientists to take the responsibility together for building, maintaining and consolidating scientific code. This seems the best solution. After all, not all scientists are programmers and vice versa. In both realms things are about to change because of the grid and service oriented architectures. In order to work in the new constellation with the grid middleware and new services and devices being brought in, significant changes will have to be made. We should leverage this opportunity, equivalent to the way the Y2K-situation led to a significant upgrade of legacy installed base in the last years of the previous millennium.

Next turn

- ➔ Develop recommendations and APIs for commercial software to interact with the European Science Grid
- ➔ Set up a support scheme on how to implement and use codes in grids and virtual environments
- ➔ Set up rules for commercial scientific software licensing to protect user communities
- ➔ Identify key open software codes and initiate a group of professional software engineers to achieve quick wins on these codes together with the user communities, and enable them to run within the grid environment

End destination

- ➔ Structural financial support for use and maintenance of scientific code, and the implementation of novel algorithms
- ➔ Modularisation and decentralised development for open and closed codes through a European software repository

Relevant policies, organisations, activities

- ➔ DG Information Society and Media, DG Research, scientific software publishers, open source developers, user groups

Directions to: Middleware repositories and parameter registration

Grids are very dynamic environments, with continuous shifts in software and middleware components, data formats and other parameters. In order to be able to replicate certain findings, one needs to make sure certain historical information is available. For data that is being used, this is done through the European storage facilities. We have also discussed software curation at the level of scientific codes. But what about the grid middleware itself? How can we preserve historical versions of middleware components and be able to find up to date versions that are compatible? How do we deal with parameters such as historical credentials or the inclusion of real time data from sensor grids?

A first requirement is that it would be very beneficial to have a central register which allocates unique parameters, name spaces, formats and schemata as used by grid applications. In scope and activities this would be quite similar to the functions that IANA performs for the internet; such data would have to be available on-line at a fixed URI in a machine-processable way (such as RDF). This bookkeeper for grid settings should probably try to operate on a global scale, and not exclude commercial parameters in order to gain global focus. This only works at the highest level, therefore it should be complemented by decentralized or local documentation of such data through semantic annotation, which will help maintain flexibility and thus may also play an important role to help organize, orchestrate or at the very least interface with services that are offered.

Another requirement is that the actual grid middleware components used in the European Science Grid remain available over a long period of time. Much of these components are open source, which greatly simplifies their archival. Several individual European countries already have comprehensive middleware repositories that actually serve a double function: since they contain all available components from historical ones until the current state of the art counterparts people can use these repositories to obtain suitable open source Grid middleware solutions. To that extent these repositories provide comprehensive information about the function, reliability and usability of every component. This is a time-consuming but important task.

Another role these middleware institutes have taken up is to provide quality-assured software engineering, testing, packaging and maintenance of software. As such they coordinate and partly finance work being done on software to make it faster and more reliable, and easy to both install and use. They also coordinate collaboration with industry. A European superstructure of middleware institutes could be an important asset for the long term, that can help avoid duplication of work, share knowledge and technology and broaden the scope of activities.

Next turn

- ➔ Open Source, production quality middleware infrastructures that unify national and European investments
- ➔ Accepted and trusted process for delivering quality, integrated software and support for large scale collaborative software development
- ➔ Support for (research into) the semantic annotation of grid services and resources

End destination

- ➔ Long term replicability of results of grid activities to ensure quality control

Relevant policies, organisations, activities

- ➔ OMII, OMEGA, VL-e, DG Information Society and Media

Directions to: Ensure open standards

The real use of the e-Infrastructures for Europe depends for a large part on the availability of mature and open protocol stacks for communication between the various components such as grids, storage components, and problem solving environments/workflow tools. In addition content and software (including virtual machines) need to be stored in a way that is interoperable and portable. Open standards²¹ are a strong enabler for sustainable long term use and development, because they enable long term access independent of any implementation and a level playing field for both non-commercial (academic, governmental, open source/free software) and all commercial players to build new infrastructures and new classes of resources and services. It is not in the interest of Europe if a commercial monopolist can hijack part of the e-infrastructure (such as grids) with proprietary exchange protocols or data formats.

In this respect working together with industry is essential, because parallel proprietary technologies – covering only a limited set of competing functionalities that allowed to develop a very dominant position outside of research and science – would be just as detrimental to the long term development of the e-infrastructure as the lack of standards. Gathering the critical mass needed for the maturing process of this technology inside and outside of science is therefore a one time opportunity, with a huge influence in the future development of the way our researchers and scientists work. Lack of standards will stifle innovation and competition. Europe should therefore invest in supporting European involvement in relevant standards bodies and standardisation efforts in a structural way, and help create a large installed base for open standards inside and outside of science and research. Standards bodies maintain the open processes that will enable growth towards a rich and mature set of standards, and not just a mandatory but incomplete set of premature recommendations which would drive people towards other technologies and achieve suboptimal results – and possibly result in a monopoly situation again. Also, in the case of ‘de iure’ standardisation (cf. ISO, ITU) such structural, vendor-independent contributions will help protect the vulnerable processes from misuse.

²¹ The following are the minimal characteristics that a specification and its attendant documents must have in order to be considered an open standard according to the clear definition by the European Interoperability Framework:

- The standard is adopted and will be maintained by a not-for-profit organisation, and its ongoing development occurs on the basis of an open decision-making procedure available to all interested parties (consensus or majority decision etc.).
- The standard has been published and the standard specification document is available either freely or at a nominal charge. It must be permissible to all to copy, distribute and use it for no fee or at a nominal fee.
- The intellectual property – i.e. patents possibly present – of (parts of) the standard is made irrevocably available on a royalty-free basis.
- There are no constraints on the re-use of the standard.

Additionally one might want to add that a standard requires the additional characteristic of its entire feature set being supported by at least two fully independent and interoperable applications.

Since much of the innovation is to come from open source software, allowing closed proprietary formats to drive parts of the grid will act as a poison pill – with tactical changes in these formats inadvertently draining development resources away from our research projects, breaking applications and thus in the end hampering innovation. Lesson learned in other ICT areas with similar traits, we should make sure that it does not happen to the grid or to other parts of the e-infrastructure.

Next turn

- ➔ Introduce travel grants to pay for more intensive European contribution to global standardisation meetings
- ➔ Appoint a European grid standards ambassador for liaising with projects and commercial stakeholders
- ➔ Ask of projects funded by EU and member states to focus and contribute to standards

End destination

- ➔ An open and mature grid protocol stack driving the global e-Infrastructure (including the European Science Grid), supporting competition and enforcing operational excellence

Relevant policies, organisations, activities

- ➔ OGF, W3C, IETF, IEEE, WS-I, IADBC, OASIS, Rosettanet, CODATA, ICTSB, NIST, DG Information Society and Media, Public Library of Science, Linux Foundation, ISO, ITU

Directions to: Training & support for scientists and support personnel

The grid environment will require many new skills for scientists on the one hand and support personnel (such as scientific software developers and academic ICT staff) on the other. Scientists need to learn how to work in new environments, conceiving and leveraging powerful new instruments. Even though user environments will try to evolve towards user-friendly interfaces, scientists will at least for the foreseeable future operate on the cutting edge of what is possible. This will no doubt spawn significant scientific rewards, but also will involve considerable effort. Developers have to learn how to write and optimise codes for use in grids so as to better utilise costly grid resources; they will also have to learn how to work within new software development frameworks. ICT staff has to be trained in supporting new applications and satisfying demanding users, while maintaining a high service level degree for the tasks they are currently already handling.

It might take years, possibly decades, before the user communities are broad and mature enough to be self-supporting, at which stage grids and problem solving environments will be a normal component of school and university curricula. The knowledge required by for instance a social scientist will be very different from the support a particle physicist will need, so disciplinary support is also essential. If we want to realise the full potential of the grid paradigm, users, developers and support personnel will need a persistent infrastructure that can provide knowledge management, education and support – both generic and geared towards specific application domains. This infrastructure may be developed and jointly exploited in cooperation with European businesses.

User support might seem to be somewhat removed from the essence of infrastructure, but it essential if people are to give up some autonomy and invest in shared facilities such as European supercomputers or storage facilities. Only if they get the same kind of high customer support they would get if they would own these resources themselves, will they accept the transfer to a European level.

In order to support self-learning, training and educational material will need to be created as open content with the proper viral licensing – so that knowledgeable users may help improve materials and help create new ones.

Next turn

- ➔ Gather, develop and maintain on-line training material for self study
- ➔ Set up technology demonstrators/training centers for scientists and others, offering knowledge and expertise both for generic technology and specific disciplinary uses
- ➔ Set up a European technical help desk (manned 24/7) for ICT support staff and developers

End destination

- ➔ Training facilities and help desk functionality for end-users, developers and support staff

Relevant policies, organisations, activities

- ➔ Grid projects, NREN's, OGF, DG Research, DG Information Society and Media, DG Education and Culture, Departments of Science/Education, ICEAGE, SELF

Resources

The European Science Grid as an integrated approach to serve the European scientific user communities should be populated with a number of resources in order for it to add value to the individual components. The word resources in this context should be interpreted in a broad way, covering literally everything that is of interest to science from computers, large storage facilities, telescopes, satellites, special physics equipment, weather balloons, lasers, spectrometers, visualisation means, and large sensor networks. A resource can also refer to large data collections, artificial intelligence agents and even people.²² The only requirement is that the resource (from supercomputer to cellphone) can at some point exchange the necessary information through standardized interfaces, i.e. grid protocols. The end goal is a rich ecosystem of resources that offer a broad gamma of hardware, software, services and data spaces.

A number of destinations are available for Resources, divided in three sections:

Building a healthy resource ecosystem

- Supercomputer Infrastructures for Europe
- European Storage Facilities
- Service Oriented instruments and facilities
- Sensor networks
- Computational accelerators

Resource Attraction and Evolution:

- Incentives for Providing Grid Resources
- Leveraging New Technologies

Directions to: Supercomputer infrastructures for Europe

Large scale computing comes in two flavours: Capability computing and Throughput or Capacity computing, describing different approaches to providing large computational power for challenging scientific applications. Capacity computing typically addresses the needs of scientific disciplines which do not need a low-latency, high-bandwidth interconnect architecture between hundreds or thousands of processors. Prices for such capacity facilities continue to drop, bringing the opportunities afforded by increased data processing and simulation supported research to a growing number of fields and problems. Capability computing, on the other hand, requires access to many processors in parallel, large memory, and low-latency, high-bandwidth interconnects capable of tackling large scale and closely coupled problems which cannot be solved in any other way. Such problems are central to progress across a wide range of scientific fields from traditional science and engineering domains to such key areas as the future of the environment, national security and public health. The boundary between capacity and capability computing is somewhat arbitrary and is not fixed, evolving as the commodity market evolves.

Within the class of capability systems, there are still trade-offs to be made. Actual architectures range from purely shared memory vector computers, through clusters of SMP systems (connected by state-of-the-art interconnects, like Infiniband) to NUMA systems, with a large amount of processors with direct access to a single address space. Since specific application areas run more efficiently on specific capability architectures, various types of systems should be part of a European Supercomputer infrastructure. Emerging technologies are quantum computing and computational accelerators, both of which can offer – for tailored applications or parts of them – several orders of magnitude increase in speed.

The need for a High Performance Computing Service for Europe integrating both capability and throughput resources has been recognized and addressed in the European Communities FP7 research program. In order to be competitive this effort will have to be increased over time and will probably remain a strong point of attention for the long term. Only through a sustained, coordinated, distributed investment in Capacity and Capability resources and expertise can Europe expect to achieve a flexible infrastructure able to respond most efficiently and effectively to the demands of the research community. An integrated infrastructure will allow researchers to exploit the technical solution most appropriate to their needs, while keeping innovation.

²² The first mass scale implementation of a computer API to steer human labour is Amazon's Mechanical Turk (M'Turk) service, introduced in 2005. The service was named after the man-powered quasi-automaton Mechanical Turk designed and constructed in 1770 by the Austrian-Hungarian baron Wolfgang von Kempelen (1734–1804). See: Gerald M. Levitt's *The Turk, Chess Automaton* (2000).

Next turn

- ➔ Set up the European High Performance Computing Service with a robust distribution of facilities across Europe, safeguarding against large scale natural and other disasters
- ➔ Invest in precompetitive procurement of Emerging Supercomputer Technologies
- ➔ Align funding policies for Research Infrastructures among European partners to enable better shared use of resources

End destination

- ➔ Continuous and secured provision of state of the art computing facilities from all important architectures available to European scientists and researchers

Relevant policies, organisations, activities

- ➔ PRACE, DG Information Society and Media, ministries of education and research, National Science Councils & Academies, ESF, Large European Computing Centres, FP7+, -

Directions to: European Storage facilities

Europe can profit from having a shared approach to the increasing storage needs and possibilities by establishing a distributed shared network of large hybrid storage facilities. Each technology by itself – e.g. in-memory storage, optical and magnetic media, holographic storage, micro arrays and biomolecular storage – has downsides for particular usage, from the point of view of cost, i/o speed, scalability, power consumption, and longevity. Therefore only a combination of technologies will be able to combine the strong points of every technology in the quest for supporting extreme speeds and huge volumes. Complementary to the supercomputers mentioned elsewhere in the roadmap some generally available gridified high-profile storage facilities will help Europe cope with the scientific data explosion and the burstiness of storage needs in large experiments in a cost-efficient and well-thought manner. This will enhance both overall performance and availability. Safe-guarding data against natural disasters (such as earthquakes and floods), technical failure, malicious intent and war will allow an uninterrupted scientific apparatus to remain operational in the most difficult circumstances. This will be all the more significant if the use of real-time simulation and predication based on data from sensor networks in such situations becomes a vital part of disaster management.

By all accounts having a shared storage infrastructure will increase peak availability of storage for the future needs of e-science. Provision of a network of distributed shared facilities will lower overall investments needed in an otherwise surging cost area as storage (unless some disruptive low-power, low cost storage technology reorders the market) by adding scale in operations and management, and taking away the need for inefficient local redundancy. The concentration of buying power and maintenance will also lower cost and increase quality, while having an installed base always ready for use lowers deployment time. It also means being able to deal better with sudden extreme surges in use of certain data or services by utilising strategies and technologies such as multicasting/anycasting, pro-active load balancing, p2p swarming (where the consumer of data becomes part of the source) and caching in large in-memory databases. This is especially relevant for real-time distributed tasks, such as responsive jobs in grids. In short, it will allow for advanced data retrieval and recovery faster than in any other scenario and at the lowest price possible – providing efficiency, flexibility, security, availability, and scalability. With the networks and grid technologies in place to provide the interconnectivity and load balancing features, shared storage facilities are a key component in the grid equation.

Next turn

- ➔ Design an optimal safe storage topology and determine a storage development roadmap
- ➔ Link large distributed storage facilities able to replicate and serve grid data as a test bed
- ➔ Find long term financial support for distributed European Storage Facilities

End destination

- ➔ A high capacity storage facility that is secure, distributed and extremely fast and capable of writing, mirroring and serving all data within the global scientific community at any given point in time

Relevant policies, organisations, activities

- ➔ e-IRG, DG Information Society and Media, National Science Councils, OGF, FP7+, ENISA, commercial storage providers (e.g. Google, Amazon S3, Ebay)

Directions to: Service Oriented instruments and facilities

Europe has a wide diversity of special scientific measuring equipment, large and small. Often these themselves have a complex internal ICT architecture. It would be very useful if these could on the one hand profit from external resources and on the other hand could themselves be accessed easily through a shared set of open standards available to the entire European Research community. Such increased availability will facilitate higher efficiency and more interdisciplinary use, and will lower the cost of these facilities themselves. To achieve such availability these resources would have to be enhanced with lightweight middleware that would make them available through services.

Prime targets are the Research Infrastructures: huge scientific instruments and installations that are globally one of a kind because of their state of the art technology, scale, or cost. Examples of existing Research Infrastructures are the Gravitational Wave Detector GEO 600 (Hannover), the ITER fusion power plant, ISIS neutron scattering facilities (near Oxford), X-ray laser XFEL (Hamburg, Schleswig-Holstein), the European Synchrotron Radiation Facility ESRF (Grenoble), Ultra Low Temperature Installation (Helsinki) and the European Molecular Biology Laboratory (Grenoble).

The RIs often involve immense investments that can only be afforded because the facilities have a life span of sometimes several decades. Due to its nature ICT provisioning for these projects over such a time line is a complex issue. Because of the unique capabilities of the RI, scientists from all over Europe (and even from across the globe) may require remote access to them. The investment in service-enabling such instruments and facilities and allowing them to use external facilities (such as HPC and extreme storage) is a first priority. In 2006 ESFRI published its first Roadmaps bringing together thirty four projects for Research Infrastructures in six categories: Social Sciences & Humanities, Environmental Sciences, Energy, Biomedical and Life Sciences, Material Sciences and Astronomy/Astrophysics/Nuclear Physics/Particle Physics.²³ These should all be service-oriented, as should Research Infrastructures that emerge in future ESFRI roadmaps.

Science is a global phenomenon. Sometimes Europe will need transparent access to scientific instruments and installations on other continents. A considerable budget goes to Research Infrastructures elsewhere on the planet. For our own cost-benefit we might consider to contribute to make such access possible on our terms on selected installations elsewhere. This would not only allow European scientists to profit better from very scarce resources not available otherwise but also would help support the strategic placement of the European approach as the standard.

Next turn

- ➔ Analyse the e-infrastructure needs of the new ESFRI projects and the e-infrastructure components of large existing projects to look for consolidation and synergy
- ➔ Initiate a funding scheme specifically dealing with service-enabling the first wave of resources, to break the impasse and create critical mass. This could be done by a Task Force that identifies and prioritizes strategic resources that should be available to the European Science community
- ➔ Install an Expertise Centre (and/or complemented with national help desks) to help deal users and owners deal with adoption issues and look into new and efficient ways to use these devices remotely and/or in a grid context

End destination

- ➔ A rich ecosystem of service-enabled devices for measuring which are able to tap into e-infrastructure resources themselves

Relevant policies, organisations, activities

- ➔ ESFRI, e-IRG, OECD, DG Information Society and Media, EUROHORCS, FP7+.

Directions to: Sensor networks

Vast amounts of data are being created all the time, through all kinds of sensors and sensor networks distributed all over the world. These involve information from a broad spectrum of application areas, from environmental sensors such as seismic data, weather data, radioactivity, electromagnetic receptors, gravitational wave observatories, pollution measurements, temperature, ground water levels and fluvial data up to measurements of human and animal activities such as (air) traffic control, critical infrastructure status or RFID-tagging of animals for zoological research. Signalling devices – such as satellites, radar equipment, radio beacons, large scale laser facilities, dikes and mobile telecom infrastructures – complement the system as they can be used to actively manipulate events that need to be measured. The Galileo European Satellite navigation system that is under construction can be used to map the virtual topology on to the real world, together with GPS and GLONASS.

Of course, many of these infrastructures are already hooked up to the research networks, but to be able to real-time interact with groups or unforeseen combinations of them, to increase the community that can use them in a sensible way, and to be able to easily integrate them into new services requires some degree of standardisation on the one hand and technological enablers such as (grid) middleware on the other. Measuring equipment is the equivalent of sensory input to our continent and the data produced sets the boundary conditions on the kind of work scientist can deliver. The combination of real-time data combined with large-scale simulation rendered through the grid will enable environment scientists to better predict what is happening and that in turn may help making better policy decisions. Especially in emergency situations, such as a vast flooding of part of Europe or a large-scale nuclear event in the middle of the continent, this might save many lives. One might also need resources that gather information about human activities on a macro scale, such as road usage, air traffic control data, electromagnetic radiation, and sound pollution. That way, Europe will gain more insight into the operational issues many of such infrastructures are facing.

Sensor data is also subject to the other needs of data with regards to storage, availability, and/or curation. Rather than having to create a redundant buffering and redistribution infrastructure for each of those continuous data outlets in order to facilitate their broad use, it would make much sense to create a universal (distributed) facility that will take care of this. Such a facility would act as a generic entry point – taking care of mirroring, time-stamping, device- and load-balancing and storage on the fly through the grid. Because of its scale it can provide a variety of access methods and APIs, and if necessary perform reliable re-mappings or translations onto other relevant data formats. This would create a multi-tier infrastructure, where primary resources only need to take care of broadcasting the data once. This would enable mobile sensors to be deployed without delay when the need arises.

Since maintaining a full copy of every bit produced by every sensor is not scalable, it would log a sensibly reduced amount of data to the European Grid Storage Facilities for long term bit preservation. In addition the reflector mechanism could in time provide an 'instant replay' buffer of for instance 48 hours that would capture full volume data on all sources to be able to provide negative latency (so that when some extraordinary event happens, one can copy the buffer and research the full data set instead of a subset).

Next turn

- ➔ Commission a reflector mechanism capable of becoming the secure and scalable front end for all European real-time data sources
- ➔ Identify and approach top 100 sensor facilities valuable to the European Science Grid to be added
- ➔ Create guidelines for conformance and provide a mechanism for all other sensors to be added at their own initiative
- ➔ Fund research in generic aggregation strategies for multiple real-time data streams and coupling with model computations and historical data

End destination

- ➔ A rich ecosystem of grid-enabled devices available for distributed measuring of all relevant environmental parameters

Relevant policies, organisations, activities

- ➔ ESA, LOFAR, ECMWF, ENBI, EMSC/IASPEI, GALILEO, and many others
- ➔ DG Information Society and Media, DG JRC, DG Environment, DG Energy and Transport, DG Agriculture, FP7+, OECD, e-Content+ programme

Directions to: Computational accelerators

An important and potentially disruptive technology trend is dedicated hardware. e-Infrastructure environments will help create ample demand for computational accelerators. These computational accelerators often take the shape of dedicated hardware, specifically designed and optimised to deal with highly specific tasks and calculations. A computational accelerator is capable of delivering extreme performance in selected tasks because the application logic is translated to a more simple wiring on the hardware level. Such hardware is known to create speed-ups of a factor 1000 or more compared to general purpose equivalents for the specific task they were designed for. Also, energy use and hardware footprint is on average significantly lower in dedicated hardware than in general purpose systems. Dedicated hardware is very stable, because it limits itself to simpler tasks only.

Dedicated hardware in the long term will play an important role in the grid as it provides blazing fast services capable of dealing with data streams reliably and cheaply at extreme speeds. Such resources clearly have many benefits to offer to e-infrastructure ecosystems, but historically designing dedicated hardware has been very time-consuming. In recent years many important developments have taken place in the field of reprogrammable logic, hardware that is either optimised for or can rewire itself for a specific task. The most popular forms are FPGAs (Field Programmable Gate Arrays), CPLDs (complex programmable logic devices) and ASICs (application-specific integrated circuits). Significant advances have been made at the manufacturing level and in the software environments used to design and craft such hardware. This has made dedicated hardware much more affordable and quicker to be developed. In time dedicated hardware may be replaced by self-configuring and self-optimising hardware.

Quantum computing is another application area that is worth serious investigation, as for selected classes of problems (including the design of more complex computers) it may offer significant opportunities. The volume and broad nature of the European grid will allow Europe to be an early adopter of these technologies and strengthen its competitiveness, but only if it makes sure that it stays upfront in the area of technology development.

In order to benefit from these technologies, research has to be done into engineering and fine-tuning of algorithms in computational accelerators.

Next turn

- ➔ Create a European centre of expertise for reprogrammable logic and computational accelerators
- ➔ Identify the most relevant and highly popular algorithms and set up a program to implement computational accelerators and quantum computers into hardware and services
- ➔ Support development of compilers and software environments that can utilise computational accelerators

End destination

- ➔ A rich ecology of dedicated hardware facilities specialized in the most time-consuming and/or most critical calculations and algorithms

Relevant policies, organisations, activities

- ➔ QIST, QUROPE, IEEE, ERCIM, ESI, ACM, DG Information Society and Media

Directions to: Incentives for providing grid resources

It is clear that there will be many different operational models for grids. Some grids will be composed of fully funded production facilities and other resources that are available for free to the communities that have access to them. Other grids will operate in market models, requiring some kind of economic return for providing specific resources. A market mechanism enables both providers of privately and publicly owned resources to invest extra in providing a broad range of services and devices at different availability and quality levels – and gives continuous incentive to maintain and update those resources.

Different negotiation models are available for grid resources, from auctions, exchanges and/or marketplaces, to agent-driven negotiators. There is no single most cost-efficient and useful mechanism, as the situation varies for various types of resources and communities. Much depends on the overhead costs of the market infrastructures for those individual communities, and the actual financial structure of the field. Therefore, the focus should be on enabling such economic models within the technical and legal domain. This means providing secure, accurate and cost-effective accounting facilities that can operate reliably within a legal framework that is supportive of the global scale of grid markets. Competing open (potentially global) exchanges and/or marketplaces for grid resources should exist. That way the organisation for allocation of resources can remain decentralised and self-organised. In order to ease the market acceptance a body overseeing Quality Assurance activities would allow users to identify different levels of guaranteed reliability for grid service providers. Security, availability, reliability and protection of privacy throughout grids are essential features for many applications of the e-Infrastructure, and these features do not follow necessarily from the mere combination of technologies. Auditability of the services and infrastructures in place is necessary to win over user communities with sensitive data or real-time requirements (e.g. medical use, financial institutions).

For resources that are being created or funded in the context of EU activities, the ability of any resource to be turned into a service available to the rest of the e-infrastructure should be an essential prerequisite for future calls. This way the incentive is built into the funding mechanism.

Next turn

- ➔ Fund work on building a functional accounting layer for grids
- ➔ Facilitate a European broker platform or entry point to rent or acquire Grid services

End destination

- ➔ A mature open market for grid resources

Relevant policies, organisations, activities

- ➔ OGE, OASIS, DG Information Society and Media, DG Trade, DG Internal Market

Directions to: Leveraging new technologies

The evolution of mainstream technology is as likely to be influenced by the introduction of e-infrastructure into the research arena as the reverse is true. Aggregated usage helps create critical mass for problems and opportunities that would otherwise stay unnoticed. A joint European e-Infrastructure will act as fertile ground for new technologies that can profit from whatever unique features these offer to deal with their rich load. Both disruptive new technologies and steady evolution of existing technologies can change the landscape radically.

In the area of interfaces and visualisation, technologies like pressure sensitive multitouch displays²⁴, volumetric displays²⁵, interactive 3D mapping of 2D media²⁶, retinal scan displays, laser plasma spatial drawing²⁷, and other state of the art technologies will conceptually alter the way people interact with their scientific data – making exploring data far more tangible and interactive. With the amount of available data growing faster than ever, the way we actually interface with it will become an important factor to success. We need to be able to send haptic and tactile information (or telekinesthetic interaction²⁸) for interaction with visualisations²⁹ but also to remotely control (robotic) devices that interact with remote environments.

The current trend of multi-core processors poses many challenges and will have a wide impact on programming models for software, including scientific software.

Next turn

- ➔ Create funding schemes that allow creation of inspiring demonstrators of emerging technologies in Europe
- ➔ Support research into haptic and tactile interfaces to facilitate virtual and remote steering
- ➔ Select a set of suitable test projects that can experiment with advanced 3D technologies

End destination

- ➔ Early access to disruptive technologies and state of the art enablers

Relevant policies, organisations, activities

- ➔ DG Information Society and Media, DG Research

24 Han, J.Y. 2005. Low-Cost Multi-Touch Sensing through Frustrated Total Internal Reflection. In: Proceedings of the 18th Annual ACM Symposium on User Interface Software and Technology

25 Kollin, J. (1993). A Retinal Display for Virtual-Environment Applications. In: Proceedings of Society for Information Display, 1993 International Symposium, Digest of Technical Papers, Vol. XXIV. (p. 827). Playa del Rey, CA: Society for Information Display.

26 Noah Snavey, Steven M. Seitz, Richard Szeliski. Photo tourism: Exploring photo collections in 3D. In ACM Transactions on Graphics (SIGGRAPH Proceedings), 25(3), 2006, 835-846.

27 Kimura, H., Uchiyama, T., and Yoshikawa, H. 2006. Laser produced 3D display in the air. In ACM SIGGRAPH 2006 Emerging Technologies (Boston, Massachusetts, July 30 - August 03, 2006). SIGGRAPH '06. ACM Press, New York, NY, 20.

28 Koike, Y., Nakakoji, K., and Yamamoto, Y. 2006. Tele-kinesthetic interaction: using hand muscles to interact with a tangible 3D object. In ACM SIGGRAPH 2006 Emerging Technologies (Boston, Massachusetts, July 30 - August 03, 2006). SIGGRAPH '06. ACM Press, New York, NY, 33.

29 The world of interactive entertainment has taken a huge lead beyond the scientific world and commercial application domain in this respect in the last decade, with clans of gamers sharing complex interactions involving photo realistic visualisation from multiple perspectives, force feedback interaction and voice/visual contact – in virtual worlds made up by millions of simultaneous users using nothing more than commodity hardware and software. For scientific uses the demands may be higher with regards to image quality, demanding better camera's and projection devices. The user interfaces will ultimately not be determined by their level of technological advancement, but by their usability. Because the human nervous system is so familiar with dealing with three-dimensional environments, 3D immersion facilities and 3D embedded projectors will start playing an important role as e-science gets more and more complex. With the current surge in 3D desktops, 3D-camera's and 3D projection technologies, the possibilities are increasing still.



Data handling

Data fuels the information age, and databases and digital libraries are the binary oil fields of that age. The amount of data is growing at an extremely high pace in what has become commonly known as the data explosion. Handling the continuous shock waves of new scientific data output and serving bursty usage patterns efficiently, while keeping that data reliable and available at the fingertips of the scientist and the historian, can only succeed with a proper data life cycle management infrastructure that leaves the distributed nature of the scientific process intact.

There is only one destination available for Data Handling:

- The Knowledge Life Cycle Infrastructure

Directions to: The Knowledge Life Cycle Infrastructure

Scientists and researchers should be able to combine and aggregate data from a multitude of sources in order to look for patterns that can lead to new knowledge - or can renew our perspective on already accepted truths. 'Existing data' is continuously used in different ways, updated, reinterpreted, corrected, forgotten, and rediscovered. The data that is interesting to the researcher might range from socio-eco-

nomical data to reference tissue sample databases, from medical records to geospatial databases, from archaeology databases to material science libraries, from genome databases to linguistic data. However, in addition to 'original' data itself, a researcher needs simple access (read and write) to metadata (data about data) that is layered on top of such knowledge, from annotations to full-scale scientific publications. The challenge is that the data sources are becoming more and more complex while the sizes of datasets will likely keep on growing by orders of magnitude. Thus, as long as we still need a human to manually link all data and metadata together, the task of overall data curation and maintenance is practically impossible to solve in a sustainable manner.

The unprecedented increase in data to be dealt with by scientists and researchers means two things: first that the (quality and availability) management of data will become even more important in the future if we are to stay in control of our data, and second that data itself is becoming a strategic resource. For the latter reason, it is of strategic importance for Europe to have data managed efficiently and to have it as close as possible. This will allow the data to be turned into services that can be used easily and sensibly via services and the grid.

By developing the data management as a service concept offered through the common e-Infrastructure standards, Europe can become the data hub of the world. It will be beneficial to Europe if it is able to guarantee transparent access to relevant data sources produced and maintained around the world. The European Storage Facilities (as proposed elsewhere in this roadmap) can be offered as a free (as in: at no cost, and no strings attached) backup and distribution mechanism for data from around the world. Such a move would be attractive for information owners and maintainers, because it provides an extra safeguard and increases general availability of their data - qualitatively and quantitatively superior than anyone may be able to afford for themselves. For Europe it is crucial that important data remains available to European scientists under any circumstance. However, a concerted effort to develop a consistent ERA-wide information ecology is necessary to fulfil these ambitions. This requires solving issues related to cost-sharing of the data curation activities and ownership of the data. The latter include various legal and technical issues, such as interoperability and "future-proofing" of DRM systems and understanding the dependencies between legacy applications and environments - both of which necessitate going beyond the ERA-area and taking a global approach.

Other benefits of Grid Storage facilities are the positive effect it may have on the 'digital divide' with scientists outside of Europe: not everyone with interesting data might be able to afford to store it adequately. Also one would expect a decreased vulnerability of the communities to politically motivated regional censorship of data from science and research. Furthermore, with supply of high-level data ser-

vices increasing, researchers and scientists will spend considerably more time in working with data beyond their own traditional domain. It becomes more affordable and attractive to provide a solid, overarching strategy for data (a data infrastructure) that can help manage the quality of and interaction between many autonomous data sources, maintainers, services and users – and contribute to technology convergence by popularising standardised, stable and user-friendly access methods to scientific data sources and digital libraries.

If data comes from many different sources, it will need to be aligned. A normalisation institute could be set up to first contribute to standardised access across organisational and international boundaries, producing validated aggregation processes and conversion schemas – in order to achieve in the long term good overall interoperability, availability and durability of scientific data. This would be complemented by support for digital libraries, semantic knowledge bases and other means to take care of data curation, software curation and semantic metadata. Without these, data loses its meaning and cannot be relied upon by scientists any more. It should be noted that in this case an organisation with processes that require human interaction would become a part of the common e-Infrastructure. Managing this in a sustainable manner requires solving issues related to sustainable funding and cost sharing models.

Next turn

- ➔ Create an enrolment mechanism for data source maintainers to use the European Grid Storage Facilities as a replicator to secure their data for free
- ➔ Identify key data sources and fully fund their addition to the European Grid storage facilities, coordinated by a Task Force that identifies and prioritizes strategic resources
- ➔ Fund research in smart replication strategies for very large databases and data sets
- ➔ Set up European repositories and digital libraries geared towards scientific software curation and serving semantic metadata
- ➔ A normalisation institute could be set up to contribute to standardised access and aggregation

End destination

- ➔ A complete and easily usable mirror (with affiliated metadata) of every significant data source in the world

Relevant policies, organisations, activities

- ➔ e-IRG, ESF, DG Information Society and Media, DG JRC, DG Eurostat, DG Internal Market, FP7+, OECD, DILIGENT, Task Force on Permanent Access to the Records of Science, European Board of National Archivists, Codata, CASPAR

Crossing the boundaries of science

In order to support Europe-wide communities that are able to interact in a global environment as equals, it is important to encourage sharing of electronic infrastructure resources as a way to create suitable conditions for cross-disciplinary interaction, providing fertile ground for innovation and eventual industrial exploitation. This has required advanced ICT collaboration tools such as sharing remote work spaces, high resolution videoconferencing, etc. These can be used in their own right, but are also very relevant to building the community around the grid even though they are generally not built on grid protocols. Collaboration and information exchange with industry – both as supplier and as a user community – and the rest of the globe is necessarily a part of the entire approach. Of course combining the major efforts from the research area and those from industry will be of great help to create a mature and sustainable market through orchestration and convergence of competing and complementary technologies.

A number of destinations are available for Crossing the boundaries of science:

- Collaboration tools and environments
- Working together with industry

Directions to: Scientific collaboration

The workbench of the scientist in the era of e-science is the scientific collaboration environment, which not only supports their work flow but enables others to contribute to that work. The scientific collaboration environment needs to be able to bring together the whole ensemble of e-infrastructure facilities scientists may require for every aspect of their tasks. The nature of the scientific process lies in argument and counter-argument, in measurements, models, and hypotheses that are in eternal uncertainty over unforeseen dependencies, error margins, relevance, and underlying assumptions. Therefore, choice and a birds eye view are the two main characteristics of healthy scientific collaboration environments. Choice, as there are many different ways to go about as well as there are vast differences among and within scientific disciplines. Birds eye view, as the daily scientific routine these days builds on the shoulders of so many giants towering on top of each other that this is bound to result in interference and conflicts to be resolved. Researchers should have the freedom to innovate or just explore their own way in each part of the scientific chain but still be able to profit from the support that technology has to offer.

Whether research communities are working together in big science environments or not, interaction within actual work processes among scientists across the globe will have to be taken significantly further than the current state of the art. We need distributed problem solving environments that offer generic and extensible frameworks to enable or support both the basic workflow (interactive visualisations, annotations, job submission and control) as well as auxiliary processes and tasks like automated software testing and debugging, logging, publication to repositories, etc. While the user environments may be highly application domain specific, it makes sense to try and build a common ground for such frameworks to interoperate – otherwise we will find ourselves stuck on non-interoperable islands in the long run.

A prime target is distributed visualisation (i.e., the possibility to look at different synchronous visualisations of the same or related data sets or streams and to simultaneously interact with this data with multiple parties at different sites and share information among them within that interaction). Of course this extends to high-grade sensory information, such as medical imagery and raw high-resolution video footage. The transformation of online human presence into the visual realm – i.e. video conferencing facilities – to enable on-line meetings is also important to take care of the human factor. Trust and mutual understanding are key elements in any collaboration and if face to face meetings are not possible these need to be mediated by virtual means.

Next turn

- ➔ Support standardisation work on workflow languages, and creation of open source software tools
- ➔ Support research and open standards for application sharing and multi-user desktop environments in collaboration environments
- ➔ Support projects funded by EU and member states in using and testing collaboration tools under development

End destination

- ➔ A set of open and mature collaboration and remote visualisation tools, compatible with the European e-infrastructure and available on all important platforms

Relevant policies, organisations, activities

- ➔ IETF, W3C, Open Group, OGF, ISO, W3C, WS-I, ICTSB, DG Information Society and Media

Directions to: Working together with industry

The e-Infrastructure will support Europe in its collaboration and competition with other global regions by allowing for better cross-fertilisation of research, technology, industrial and commercial services and products within the region itself and with the outside world. It will help achieve significant upscaling of technological possibilities through virtualisation and increased specialisation in hardware and skills. Getting the conditions right for cross-disciplinary interaction is essential for sustained competitiveness of Europe on a global scale. With Europe covering less than 10 percent of the global population, collaboration with the rest of the globe is necessarily also a part of the entire approach.

The academic, private, and public domains have strong autonomous drivers for their respective developments. Nurturing the relationships between these domains is therefore very important, as each domain may act as innovator and supplier on the one hand and as a large user community on the other. Combining the major efforts from the research area and those from industry will be of great help in creating a mature and sustainable market through orchestration and convergence of competing and complementary technologies. Such collaboration is non-trivial, and takes vision to tailor policies, acquisition strategy and other incentives. It also takes serious investments in building and maintaining relationships.

Although grid technology for instance may have its roots in science, the concept translates itself easily towards industry and promises to be of considerable economic influence. As was witnessed with the development of the internet and the world wide web, the influx of a large volume of commercial and governmental users greatly enlarges the possibilities if the efforts are combined. Science, industry, and the public sector should work together in order to make sure that a set of open standards and a broad community of services that is of use to all emerges. Only through combined volume of activities will we gain critical mass and achieve interoperability. There is a joint group interest in creating a mature future for grids that outside artificial monopolies based on proprietary grid-like standards. Rather than directly interfacing with individual market players from every corner of the industry the focus should be on getting the boundary conditions right and giving the right incentives. Creating a dependability on 'local champions' with proprietary technologies is very fragile in a global economy where instant shifts of ownership and power are just a matter of buying stocks.

When we think of the ICT environments of the future, surely the developments within the media and entertainment sector play an important role. Although these developments are seemingly at a distance from the scope of the strict idea of e-Infrastructure, the investments in their value chain – from real estate owners, network operators, content creators and entertainment corporations – will ultimately have a major impact on the way the e-Infrastructure will develop. The same goes for generic ICT spending in the private sector and by governments.³⁰

Next turn

- ➔ Invite industry to participate in the further development of the Roadmap for e-Infrastructures

End destination

- ➔ Interoperable and advanced grid (and other ICT technologies) based on shared standards used by science and industry

Relevant policies, organisations, activities

- ➔ e-IRG, IDABC, DG Information Society and Media, DG Internal Market, OGF, OASIS, Rosettanet, European standards bodies

³⁰ To illustrate this: the successors of the cheap computer chips for the business and home market (for desktop tasks and gaming) have almost completely overtaken the multi-billion euro market share of the special purpose design chips in high tech supercomputers for science and research, because the market volume of these commodity processor chips is so huge compared to the latter. This has had direct consequences for the software environments that could be run on supercomputers and impacted scientists using them immediately - depending on how badly they relied on specific features from the old architecture.

Abbreviations and acronyms

CERN	European Organization for Nuclear Research, the world's largest particle physics laboratory located near Geneva, Switzerland.	www.cern.ch
DANTE	Delivery of Advanced Network Technology to Europe.	www.dante.net
DF	Dark Fiber, i.e. fiber optic cables physically present between locations but not yet used by the owner – and therefore available to be leased or sold to others.	
DILIGENT	Digital Library Infrastructure on Grid ENabled Technology, an EU FP6 project on digital libraries.	diligentproject.org
e-Content+ programme	Programme within the EC Sixth Framework Programme to stimulate the development and use of European digital content.	
e-IRG	e-Infrastructures Reflection Group, a policy body consisting of national delegates that defines and recommends best practices for each of the (pan-)European e-Infrastructure efforts.	www.e-irg.eu
ECMWF	European Centre for Medium-Range Weather Forecasts. an international intergovernmental organisation for weather prediction based in England.	www.ecmwf.int
Eduroam	Eduroam stands for Education Roaming and is a RADIUS-based infrastructure to allow inter-institutional roaming.	www.eduroam.org
EMSC	European-Mediterranean Seismological Centre.	www.emsc-csem.org
ENBI	European Network for Biodiversity Information, a thematic network aims at coordinating Europe's efforts in the broad field of biodiversity information.	www.enbi.info
ENISA	European Network and Information Security Agency, is an agency of the E established to improve network and information security in the European Union.	www.enisa.eu.int
ESA	European Space Agency (ESA) is an inter-governmental organisation dedicated to the exploration of space.	www.esa.int
ESF	European Science Foundation.	www.esf.org
ESFRI	European Strategic Forum for Research Infrastructures.	cordis.europa.eu/esfri
EugridPMA	European Policy Management Authority for Grid Authentication.	www.eugridpma.org
EuroPKI	A not-for-profit organization established to create and develop a pan-european public-key infrastructure (PKI).	www.europki.org
FP7+	Seventh Framework Programme, the upcoming (2007-2013) Framework Programme for Research and Technological Development set up by the EU.	ec.europa.eu/research
FPGA	Field Programmable Gate Array, a semiconductor device containing reprogrammable logic and interconnects.	
GALILEO	A satellite navigation system independent from GPS and GLO-NASS that has been commissioned by the European Union.	ec.europa.eu/dgs/energy_transport/galileo

GEANT	The pan-European research network, currently in its second incarnation, GEANT2.	www.geant2.net
GLIF	Global Lambda Integrated Facility. An international virtual organization that promotes the paradigm of lambda networking.	www.glif.is
GLORIAD	Global Ring Network for Advanced Applications Development.	www.gloriad.org
IANA	Internet Assigned Numbers Authority. IANA controls numbers for protocols, the Country Code Top Level Domains and maintains the IP Address allotments.	www.iana.org
IASPEI	International Association of Seismology and Physics of the Earth's Interior.	www.iaspei.org
ICTSB	ICT Standards Board, a European coordination platform for specification activities in the field of Information and Communications Technologies (ICT) initiated by the three recognized European standards organizations CEN, CENELEC and ETSI.	www.ictsb.org
IDABC	Interoperable Delivery of European eGovernment Services to public Administrations, Businesses and Citizens, an EU programme aimed at improving efficiency and collaboration between European public administrations.	europa.eu.int/idabc
IEEE	Institute of Electrical and Electronics Engineers, a global professional engineering association responsible for many electrotechnical standards.	www.ieee.org
IETF	Internet Engineering Task Force, the standardisation body that establishes the internet standards.	www.ietf.org
IP	Internet Protocol.	
IPv6	Internet Protocol version 6, the latest version of the protocol that runs the internet.	www3.ietf.org/rfc/rfc2460.txt
ITU	International Telecommunications Union, an international organization established to standardise and regulate international radio and telecommunications.	www.itu.int
LAN	Local Area Network, a small scale computer network.	
LOFAR	LOW Frequency ARray for radio astronomy, a large sensor grid based in The Netherlands.	www.lofar.org
NREN	National Research & Educational Network, the national entity responsible for providing network access and services to the research and education community.	
NRENPC	Policy Committee within GEANT2 with appointed representatives from each partner in the project.	www.geant2.net
NUMA	Non-Uniform Memory Access.	
OASIS	Organization for the Advancement of Structured Information Standards, a global consortium that develops and supports convergence and adoption of e-business and web service standards.	www.oasis-open.org

OECD	Organisation for Economic Co-operation and Development, an international organisation that provides collaboration on policy issues.	www.oecd.org
OGF	Open Grid Forum, the result of a merger between the Global Grid Forum and the Enterprise Grid Alliance.	www.ogf.org
OMEGA	Open Middleware Enabling Grid Applications.	
OMII	Open Middleware Infrastructure Institute, a repository of interoperable and open-source Grid middleware established within the UK e-Science Programme.	www.omii.ac.uk
Open Group	The Open Group is an industry consortium to set vendor- and technology-neutral open standards for computing infrastructure.	www.opengroup.org
RDF	Resource Description Framework, a semantic technology for metadata.	www.w3.org/RDF
RFID	Radio Frequency Identification, a tagging method based on remotely readable electromagnetic devices.	
RI	Research Infrastructures.	
SEEREN	South-East European Research and Education Networking project.	www.seeren.org
SMP	Symmetric multiprocessing, the use of multiple CPUs.	
TERENA/TACAR	Trans-European Research and Education Networking Association is an association of organisations that are involved with the provision and use of computer network infrastructure and services for research and education in Europe.	www.terena.nl
URI	Uniform Resource Identifiers, a way to assign a unique address to an object in an information space.	www.w3.org/Addressing
VL-e	Virtual Laboratory for e-Science, a project building e-Science tools.	www.vl-e.nl
VoIP	Voice over IP, the ability to have audio conversations over an IP network.	
W3C	World Wide Web Consortium, an international consortium developing the standards for the World Wide Web.	www.w3.org
WAN	Wide Area Network, a computer network covering a wide geographic area.	
WS-I	Web Services Interoperability Organization, an industry effort that promotes Web Services Interoperability.	www.ws-i.org

e-IRG
e-Infrastructure
Reflection Group



e-IRG
e-Infrastructure
Reflection Group

<http://www.e-irg.eu>

e-Infrastructures Roadmap 2007

