e-IRG

e-Infrastructure Reflection Group

e-Infrastructures Roadmap

e-IRG

e-Infrastructure Reflection Group

http://www.e-irg.eu

Published by e-IRGSP

e-Infrastructures Roadmap

Leenaars, M.A.G.J., Karayannis, F. et al (ed.)
First edition, version approved by e-IRG on december 13. 2005.

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.

e-Infrastructures Roadmap

About this document

The term e-Infrastructure indicates the new generation of ICT-based Research Infrastructure in Europe. The purpose of the e-Infrastructure Roadmap is to outline the necessary steps Europe should take in the field of e-Infrastructure to realise the vision of the i2010 agenda and beyond. With a time-line of twenty years we will list the most promising opportunities in this area for coordinated European action from the member states and targeted EC funding from both the strategic and the cost/benefit point of view. The document was produced by the e-Infrastructure Reflection Group and is to be updated on a regular basis in order to incorporate future developments. The latest version can always be found at www.e-IRG.eu/roadmap.

Overview of contributors

e-IRG delegates:

Werner Kobliz	*(Aconet Bundesm. für Bildung, Wissenschaft und Kultur) Austria*
Dieter Kranzlmueller	*(GUP/Uni Linz) Austria*
Pierre Bruyere	*(Belnet) Belgium*
Rosette Vandenbroucke	*(Vrije Universiteit Brussel) Belgium*
Mario Campolargo	*(European Commission/DG Information Society) EU*
Kyriakos Baxevanidis	*(European Commission/DG Information Society) EU*
Jan Gruntorad	*(Cesnet) Czech Republic*
Jan Windmuller	*(Ministry of Science Technology and Innovation) Denmark*
Jaak Anton	*(ICT advisor, Ministry of Education and Research) Estonia*
Indrek Reimand	*(Head of Research Division) Estonia*
Leif Laaksonen	*(CSC) Finland*
Dany Vandromme	*(RENATER) France*
Ruediger Krahl	*(Deutsches Zentrum fuer Luft- und Raumfahrt e.V.) Germany*
Wolfgang Gentzsch	*(BMBF/D-Grid) Germany*
Helmut Löwe	*(Bundesministerium für Bildung und Forschung) Germany*
Fotis Karayannis	*(GRNET / Hellasgrid) Greece*
Panayiotis Tsanakas	*(GRNET / Hellasgrid) Greece*
Peter Hanak	*(National Office for Science and Technology) Hungary*
Lagos Bálint	*(NIIF/Hungarnet) Hungary*
Richard Hirsh	*(SFI) Ireland*
Tom Sheedy	*(International Programmes, Enterprise Ireland) Ireland*
Mirco Mazzucato	*(Istituto Nazionale di Fisica Nucleare) Italy*
Francesco Beltrame	*(MIUR) Italy*
Ilmars Slaidins	*(Riga Technical University) Latvia*
Laimutis Telksnys	*(LITNET/Institute of Mathematics and Informatics) Lithuania*
Theo Duhautpas	*(RESTENA Foundation) Luxembourg*
Antoine Barthel	*(RESTENA Foundation) Luxembourg*
Robert Sultana	*(University of Malta) Malta*
Patrick Aerts	*(NCF/NWO) Netherlands*
Kors Bos	*(NIKHEF) Netherlands*
Gudmund Høst	*(The Research Council of Norway) Norway*
Rolf Larsen	*(Ministry of Education and Research, Dept. of Higher Education) Norway*
Norbert Meyer	*(Poznan Supercomputing and Networking Center) Poland*
Robert Pasínski	*(Ministry of Scientific Research and Information Technology) Poland*
Gaspar Barreira	*(Instituto de Prospectiva) Portugal*
Dorin Carstoiu	*(Politehnica University of Bucharest) Romania*
Branko Zebec	*(Ministry of Education Science and Sport) Slovenia*
Sergi Girona	*(Barcelona Supercomputing Center) Spain*
Luiz E. Ruiz	*(Ministry of Science and Technology) Spain*
Torbjörn Wiberg	*(ICT-Unit, Umea University) Sweden*
Anders Ynnerman	*(SNIC) Sweden*
Fiorenzo Scaroni	*(SWITCH) Switzerland*
Malcolm Read	*(JISC) UK*
Anthony Hey	*(Engineering & Physical Sciences Research Council) UK*
Alistair Dunlop	*(Open Middleware Infrastructure Institute) UK*

e-IRG observers:

Victor Alessandrini	*(IDRIS – CNRS) France*
Vladimir Getov	*(Univ. of Westminster and Bulgarian Academy of Science) UK*
Bob Jones	*(CERN) Switzerland*
Gagliardi Fabrizio	*(CERN) Switzerland*
Vasilis Maglaris	*(NREN PC) Greece*

Expert input to the Roadmap:

Borka Jerman Blazic	*(ISOC-ECC)*
Ioan Dumitrache	*(NURC)*
Axel Berg	*(SARA) Netherlands*
Jorge Gomez	*(LIP)*
David Groep	*(NIKHEF, EUgridPMA)*
Petter Kongshaug	*(UNINETT)*
Jacko Koster	*(University of Bergen)*
Mikhel Kraav	*(EENet)*
Rossend Llurba	*(NCF)*
Diego Lopez	*(RedIRIS)*
Peter Michielse	*(NCF)*
Michel Spiro	*(IN2P3/CNRS)*
Selçuk Taral	*(UBITAK)*
Ziga Turk	*(University of Ljubljana)*

Editorial team:

Michiel Leenaars	*(e-IRGSP) -lead editor*
Matti Heikkurinen	*(e-IRGSP)*
Panagiotis Louridas	*(e-IRGSP)*
Fotis Karayannis	*(e-IRGSP)*

Additional editorial support:

Naomi Messing	*(e-IRGSP)*
Hanneke van Kempen	*(Vrije Universiteit)*

Table of contents

Introduction

In the 20th century Europeans produced the first computer,[1] invented packet switching[2] (the basis for the technology operating the internet) and more recently conceived the *world wide web*.[3] 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 greenfield 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, 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. This will provide fertile ground for innovation and eventual industrial exploitation and use in education.

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

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 technologies.[4]

What is an e-Infrastructure?

The term e-Infrastructure is used to indicate the integrated 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 to 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 numerically 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 physics and high-energy physics.[5] 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 an strong trend towards mass deployment of ICT to manage the large variety of decentralised data sources and find novel approaches to traditional problems.

4 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.
5 For more examples see: 'Recommendation on the Installation of European Supercomputers', Wissenschafsrat, november 2004.

Key components of the e-Infrastructure are *networking infrastructures, middleware and organisation* and various types of *resources* (such as supercomputers, sensors, data and storage facilities). In diagram 1.1, the relationship between those components becomes clear. The network is at the heart of everything. In the age of hybrid networks, 'internet' should probably just read 'network' in order to also apply to lambda-networking. The middleware and virtual organisations (exemplified by the black knotted triangle) connect the distributed resources, data and storage facilities in a seamless way. 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.

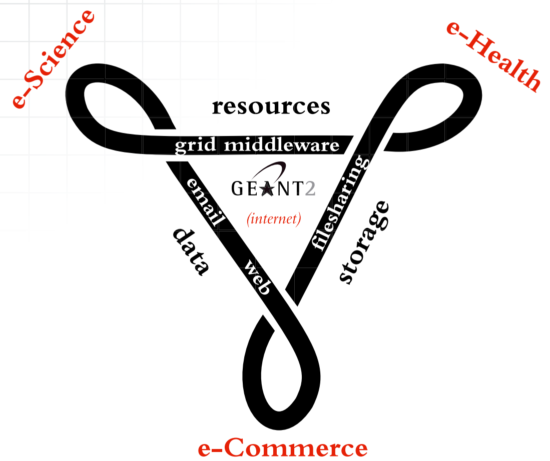

1.1 A schematic overview of e-Infrastructure components.

Let's take a closer look at each of those components:

- **Networking Infrastructure**
  The research networking infrastructure delivers the physical connections for the e-Infrastructure. These will primarily be delivered through the hybrid GÉANTx[6] 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 GÉANTx 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 the shift from discipline oriented solutions to generic facilities 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 large data collections, 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.

Grid middleware is complemented by other advanced ICT collaboration tools – such as sharing remote work spaces, high resolution videoconferencing. 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.

6 Where GÉANTx denotes any future incarnation of the GÉANT/GÉANT2 network.

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 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 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 NREN's 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 system currently in place) is worth over 0.6 billion Euro.[7] 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 is spent on large Research Infrastructures,[8] 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.[9]

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.[10] 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.[11]

7 Estimate derived from the ARCADE-EU database, www.arcade-eu.org.
8 See: www.cordis.lu.
9 See: www.lofar.org.
10 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.
11 In the case of supercomputing facilities commercial offerings from companies like IBM and SUN 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. Let's take the example of supercomputers: during the last twenty years it took 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.[12] 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 pharmaceutics 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.[13] 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 man 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.[13] 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.[15] 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.

12 Estimate: Netherlands National Computing facilities foundation, 2005.
13 Source: Centre for Medicines Research International, www.cmr.com.
14 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.
15 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.

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 ultimately in the world) 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 won't 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 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.

Additional advantages

There are many more good reasons why combining resources scattered across the continent is a key part of our European future. Besides the advantages mentioned before, adopting the grid paradigm will also yield the following advantages or new options compared to non-networked resources:

1. Connecting resources that necessarily are in a different place or are owned and maintained by different parties.
2. Combining investments to lower price and/or to enhance quality, volume or achieve certain special capabilities and distribute these resources fairly among the partners.
3. Leveraging what would otherwise be written off as temporary local overcapacity – either selling, bartering or giving it away.
4. Allowing a fully decentralised set-up to profit from the benefits of centralisation without (many of the) penalties. An example is flexible just in time federation of information coming from many data sources being maintained by many independent (even competing) parties.
5. Widening of available systems. Diversity both in type and age of the available machines on the grid is actually an asset, e.g. for legacy (versions of) applications that are dependent on certain hardware or software platforms.
6. Dynamic load balancing and data replication to increase performance, reliability and resilience.
7. Once a default method to connect facilities to the e-Infrastructure is set, contributing additional resources suitable for more general use is simple and straightforward. People can directly profit from joining by instantaneous interoperability with a very large pool of other resources.
8. Building a community around those combined resources that is able to cross-fertilise easily because technical and organisational barriers have been lifted.

How to read this Roadmap

Networking, middleware and resources as described above together make up the e-Infrastructure. In the roadmap we will define for all these components from the idea of Europe as one inter-connected 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 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

- Centrally owned or largegrained, shared use:
  - Supercomputer infrastructures for Europe
  - European Storage Facilities
  - Making sense of sensors
- Data for the grid (Locally owned, pooled and shared)
  - Grid-enabled instruments
  - Leveraging new technologies
- Incentives for providing grid resources

Crossing the boundaries of science

- Collaboration tools and environments
- Working together with industry

It will be challenging to make this agenda come through. The developments in many of these areas are still going at an extreme pace with increasingly higher stakes. 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 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 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

A next generation optical pan-European network platform GÉANT2 was launched in 2005. This will integrate 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

Directions to: 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 heavily invested 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 (DF) 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 cost effectiveness sake, building dedicated network of links between major computing centres would be a good investment. This is due to the very advantageous cost/performance ratio of so called 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 doorstep of households, companies, governments, institutes and organisations is a key factor of the development of the e-Infrastructure in the long term 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. A European organisation should be founded that investigates these opportunities and if necessary subsequently starts up 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:

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

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 GÉANTx 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 the shift from discipline oriented solutions to generic facilities special attention is needed to support and train users.

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 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 (SME's) 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 ID card* 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:

- ➡ 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, International Grid Trust Federation.

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 its implements incremental knowledge, approaches, algorithms and models available. 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. The software is in some areas – where complexity has risen to the point that researchers cannot but depend on the validity and accuracy of the code – 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 maintaining their pace of publication, financially draining the community to the point of extortion is lying around the corner. Current Intellectual Property Right solutions are not in the interest of science. The only exit-scenario is to rebuild from scratch as a community effort, which is very time-consuming as these codes have been worked on for many man years and the internals of the software are unclear to outsiders. 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 peoples 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 aren't 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 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 would 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 have professional software engineers and scientists take the responsibility together for building, maintaining and consolidating scientific code. This seems the best solution. After all, not all scientists are programmers.

In both realms things are about to change because of the grid. In order to work in the new constellation with the grid middleware and new devices being brought in, significant changes will have to be made. We should leverage this, equivalent to the leveraging of the Y2K-situation that led to a significant upgrade of legacy installed base in the last years of the previous millennium.

Next turn

- ➡ Develop recommendations for commercial software to interact with the European Science Grid.
- ➡ Set up a support scheme on how to implement and use codes in grids.
- ➡ Develop adequate regulatory framework for commercial software licensing to protect user communities.
- ➡ Identify the most important 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 schemas 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 per se 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 least interface with services that are offered.

One other 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 archiving.[16] 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.

16 Since people sometimes forget to share or document any changes to the code they make themselves – thereby using the open source model asymmetrically – such a middleware repository should probably try to at least capture a binary version of whatever codes are run if they are used for jobs of a certain size.

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 use of e-Infrastructures for Europe depends for a large part on the availability of a mature and open grid protocol stack. Open standards are a strong enabler for sustainable long term use and development, because they enable a level playing field for both non-commercial (academic, governmental, *open source*) and all commercial players to build new infrastructures and new classes of resources and services. Europe does not want a commercial monopolist hijacking the grid with proprietary exchange protocols - or in command of a parallel technology covering a limited set of competing functionalities, as this would be just as detrimental to the critical mass needed for the maturing process of grid technology. Lack of standards will stifle innovation and competition. It should therefore strongly support European involvement in the relevant standards bodies in a structural way, and help create a large installed base for open standards.The 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 end up in a monopoly situation again there.

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 bleeding competition to death.As this has recently happened in many ICT areas, we should make sure that it does not happen to the grid.

Next turn

- Introduce travel grants to pay for more intensive European contribution to global standardisation meetings.
- Appoint a European grid standards ambassador to liaison 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 European Science Grid, supporting competition and enforcing operational excellence.

Relevant policies, organisations, activities:

- OGF,W3C, IETF, IEEE,WS-I, ICTSB, DG Information Society and Media.

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 social scientists will be very different from the support particle physicists 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.

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 special 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, DG Research, DG Information Society and Media, DG Education and Culture, Departments of Science, Education and Culture.

Directions to: Incentives for providing grid resources

It is clear that there will be many different operational models for grids. Some grids will exist 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 owned resource 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.

Various 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, this 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 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 integrate into a grid (i.e. for resources to be grid-ready from the start) 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:

- ➡ OGF, OASIS, DG Information Society and Media, DG Trade, DG Internal Market.

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 as support organizations that can be shared between institutes. The only requirement is that the resource (from supercomputer to cell phone) 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:

- Centrally owned or large grained, shared use:
  - Supercomputer infrastructures for Europe
  - European Storage Facilities
  - Making sense of sensors
- Data for the grid (Locally owned, pooled and shared)
  - Grid-enabled instruments
  - Leveraging new technologies
  - Incentives for providing grid resources

Directions to: Supercomputer infrastructures for Europe

Large computing facilities come in two flavours: capacity computing and capability computing. Capacity computing typically fills the needs of scientific disciplines which do not need a low-latency, high-bandwidth interconnect architecture between processors. Prices for such capacity facilities continue to drop, and it is not expected that combining these national resources into a few capacity super clusters in Europe will benefit from the economy of scale. Surely there exists a need for large cluster facilities, but these tend to fall well within the financial possibilities of users themselves.

Capability computing, on the other hand, needs a European approach. In general capacity computing requires access to many processors in parallel, large memory, and low-latency, high-bandwidth interconnects. The availability of significantly larger capability systems than available now is urgently required for increasing the scientific opportunities for many application areas, since more processors and memory will be available in a low-latency, high-bandwidth interconnect architecture. Hence, the availability of a number of capability systems will facilitate economically important fields of science and research in Europe in their efforts of running increasingly accurate simulations. 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, it may be desirable to represent these types of systems in a European Supercomputer infrastructure.

Next turn

- ➡ Perform gap analysis on missing classes of supercomputers architectures.
- ➡ Design a robust distribution of facilities across Europe, safeguarding against large scale (natural)disasters.
- ➡ Send out a Request for Proposals for the provision of shared European supercomputers.
- ➡ 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.

Relevant policies, organisations, activities:

- ➡ 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 storage facilities. Complementary to the supercomputers and the large Research Infrastructures some gridified high-profile storage facilities will help Europe cope with the scientific data explosion in a cost-efficient and well thought manner. This will enhance both performance and availability. Safe-guarding data against natural disasters – such as tsunamis, earth quakes, floods – technical failure and malintent will allow an undisrupted 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 exceptional uses science will (almost by definition) come up with. Provision of a network of distributed shared facilities will reduce overall costs as it takes 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 ready for use any time lowers deployment time. Grids are able to deal with sudden popularity of data, using the swarming effect (the consumer of data becomes part of the source). In short, it will allow for advanced data 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 European Grid storage facility that is secure, distributed and extremely fast. This high capacity storage facility is at any given point in time capable of mirroring and serving all data within the global scientific community.

Relevant policies, organisations, activities:

- ➡ e-IRG, DG Information Society and Media, National Science Councils, OGF, FP7+, ENISA.

Directions to: Data for the grid

The European Grid Storage Facilities can be used to replicate and provide standardised access to scientific data sources and digital libraries. These range from socio-economical data to reference tissue sample databases, from astronomical records to geospatial databases, from archeology databases to chemistry libraries, from genome databases to linguistic data. Data fuels the information age, and databases and digital libraries are the binary oil fields of that age. Scientists 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. This involves data from as many different fields of science, arts and humanities as possible, as well as relevant data from other (such as commercial) sources.

It will be beneficial to Europe if it can have transparent access to relevant data sources produced and maintained around the world. The distributed storage facilities in the European Science grid can be offered as a free replicator mechanism for data from around the world. Such a move would be attractive for information owners and maintainers, because it provides an extra safe-guard and increases general availability of their data. For Europe it is important the data remains available to European scientists under any circumstance.

It is of strategic importance to Europe to have data as close as possible and develop the data into services that can be used easily and sensibly across the grid – and beyond. Europe can in this way become the data hub of the world – the heart chamber of the information age - but a concerted effort to develop a consistent internal information ecology is necessary to fulfil these ambitions.

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 and other means to take care of data curation, software curation and semantic metadata. Without these, data loses its meaning and cannot be transferred to knowledge by scientists any more.

Next turn:

- ➡ Create an enrolment mechanism for data source maintainers to use the European Grid storage facilities as a replicator to secure at least one copy of 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 replication strategies for very large database.
- ➡ 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, available either real-time or with a time lag.

Relevant policies, organisations, activities:

- ➡ e-IRG, ESF, DG Information Society and Media, DG JRC, DG Eurostat, DG Internal Market, FP7+, OECD, DILIGENT.

Directions to: Making sense of sensors

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. Signalling devices – such as satellites, radar equipment, radio beakons, large scale laser facilities 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 about to be built can be used to map the virtual topology on to the real world.

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 unexpected 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 one hand and technological enablers such as grid middleware on the other. Measuring equipment is the equivalent of sensory input of our continent and the data they produce together sets some 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 scientists to better predict what is happening and help make better policy decisions. Especially in emergency situations, such as a vast flooding of part of Europe or a large-scale nuclear accident 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 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 redundantly create a 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 fuse for their access – taking care of load balancing on the fly through the grid. Also, it can provide 'translations' into a diversity of protocols. 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 preservation. 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 a.s.a.p.
- ➡ 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: Grid-enabled instruments

Europe has a wide diversity of special scientific measuring equipment, large and small. It would be very useful if these could 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. To achieve such availability these resources would have to be enhanced with lightweight middleware that would make them available within grid infrastructures.

'Prime target are the Research Infrastructures (RI)', a name used for huge scientific instruments and installations that are globally one of a kind because of their state of the art technology, scale or cost. Examples 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), European Molecular Biology Laboratory (Grenoble) and the European Spallation Source (to be allocated). The RI often involve immense investments that can only be justified because they have a life span of sometimes several decades. Due to the unique nature of the RI, scientists from all over Europe (and even from across the globe) may require remote access to them. The investment in grid-enabling such instruments is a first priority, and will be worth every Euro. ESFRI will in its roadmaps that are developed parallel to this roadmap concern itself with trying to structure future Research Infrastructures, which will need to also be grid-enabled.

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 picking up the bill for the investment 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:

➡ Initiate a funding scheme specifically dealing with gridifying 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 Grid.

➡ Install an Expertise Center (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 gridify these devices.

End destination

➡ A rich ecosystem of grid-enabled devices for measuring.

Relevant policies, organisations, activities:

➡ ESFRI, e-IRG, OECD, DG Information Society and Media, CERN, FP7+.

Directions to: Leveraging new technologies

The introduction of grids into the research arena creates opportunities that were hitherto unavailable, which will influence the evolution of technology. Because the volume of use of a grid can be much much larger than that of any regular system, new technologies can profit from whatever unique features they offer to deal with this rich load.

One such technology is reprogrammable logic. Grid environments will create critical mass for dedicated hardware, specifically designed and optimised to deal with highly specific tasks (such as digital signal processing) or calculations. Dedicated hardware should be seen as targeted coprocessors, capable of delivering extreme performance in selected tasks. Because the application logic is translated to a hardware level, such hardware can create speed-ups of a factor 1000 or more – for specific tasks only – compared to general purpose equivalents. 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 one task 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 grid ecosystems, but historically designing dedicated hardware has been very time-consuming. In recent years many important developments have for instance been made in the field of FPGA's (Field Programmable Gateway Array's) and the software environments used to design and create such hardware. This has made dedicated hardware affordable, just in time for the grid era.

In general, new directions for computer architectures such as utilizing multi-core processors should be monitored for their applicability in the European Science Grid. The volume and broad nature of the European grid will allow Europe to be early adopters for these technologies and strengthen its competitiveness.

Next turn

- ➡ Identify relevant and highly popular tasks which could profit from new directions in hardware and software such as implementation in reprogrammable logic.
- ➡ Create a European center of expertise for reprogrammable logic and dedicated hardware.

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:

- ➡ DG Information Society and Media, DG Research.

Crossing the boundaries of science

In order to support European 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 requires advanced ICT collaboration tools such as sharing remote work spaces, high resolution videoconferencing. 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: Collaboration tools and environments

If the European Research Area is to take off, serious investments need to be made in order to develop and provide better collaboration environments. Whether research communities are working together in grid environments or not, sharing information across Europe and beyond will have to be taken significantly further. In the nineteen seventies the first ideas for groupware as we now call it were formed in the visionary Plato project, but not much has happened since. Most groupware is still (out of historical grounds) focussed on providing a very basic set of tools for a closed LAN or WAN-environment, although now of course groupware solutions extend their functionality to communicate also to the outside world. Many of the classical groupware features such as shared storage, calendaring, e-mail/messaging and a forum/discussion area are no longer unique to groupware – and often found in superior forms on their own.[17] Although a lack of good standards can be blamed for keeping better solutions from being developed, the problem until now was probably also that for new conceptual ground to be broken in collaboration environments the networks, the hardware and software were not ready.

If we assume the software and hardware now are (or will be within the time frame we are now looking at), what modern collaboration tools do we need? We can take e-mail, calendaring, static file-sharing, VoIP, wiki's/web fora and secure instant messaging for granted, although embedding these in the user environment would be nice, but otherwise these tools come 'for free' with the rest of the e-Infrastructure. First of al we need to share run-time software and interaction data between clients: this can be done either through a client-server/p2p model or through transmitting the application itself and running it on the local machine(s) while sending update messages between the various participants. Bulk data can be provided through the grid. Such a set-up will give us 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 – and share information about this within that interaction). It will also enable complex distributed problem solving environments. The framework for these should be generic, while the user environment will be highly application domain specific.

17 Probably the largest advanced international ICT-based collaborations are the one within the open source domain, such as the development of the Linux kernel and the Debian developers. These open collaborative communities share many characteristics with the science domain, yet they hardly use tools from the classical groupware solutions. They came into being by using what was available to them: generic internet services such as Usenet, e-mail and instant messaging as their communication platform, with FTP-servers and web-servers to deal with the distribution of their work. Distributed development on the lowest level is done through tools such as CVS and Subversion for code and wiki's for written text. Many people from the academic community contribute to open source projects, and find this toolset to be highly efficient. Of course, these are highly skilled users with a strong ICT profile; probably the average scientist will need some help in initially setting up these services.

Secondly, we need to to share high-grade sensory information, such as medical imagery and raw high-resolution video footage. In addition we need to be able to send haptic and tactile information in order to be able to remotely control (robotic) devices that interact with the remote environment. 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. Also, the interaction with services through the grid will by its nature be much more diverse. 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-camera's and 3D projection technologies, the possibilities are increasing still.

Of course video conferencing facilities to enable on-line meetings are also important to take care of the human factor: trust and mutual understanding are key elements in any collaboration. Face to face meetings, mediated by virtual means, are an important component.

Next turn

➡ Support research and open standards for application sharing and multi-user desktop environments in collaboration environments.
➡ Support research into haptic and tactile interfaces to facilitate remote steering.
➡ Select a set of suitable test projects that can experiment with advanced 3D camera's and displays.
➡ 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 Science Grid and available on all important platforms.

Relevant policies, organisations, activities:

➡ IETF, Open Group, OGF, W3C, WS-I, ICTSB, VL-e, DG Information Society and Media.

Directions to: Working together with industry

Grid technology 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. Perhaps it is equally interesting to look at areas such as rich collaboration environments, which have also large predicted uses but have so far more or less failed to develop.

Science and industry should work closely 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. Of course industry and science are heterogeneous categories, but there is a joint group interest in creating a mature future for grids that bans out artificial monopolies based on proprietary grid-like non-standardised protocols, API's and file formats.

Next turn

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

End destination

➡ Interoperable and advanced grids based on shared standards used by science and industry.

Relevant policies, organisations, activities:

➡ OGF, e-IRG, DG Information Society and Media, standards bodies, DG Internal Market.

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 predication 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 touropean Union 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

GÉANT	The pan-European research network, currently in its second incarnation, GÉANT2.	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 plagform 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