

European Grid Infrastructure

2011–12



Contents

Introduction	1
What does EGI do?	3
EGI's European impact.	5
Supporting the EGI ecosystem.	12
EGI's strategy for the future	21
European projects	24
EGI.eu financial statement 2011	28
EGI.eu management structure	30
Glossary.	32

Why e-infrastructures?

EGI: an e-infrastructure for the European Research Area

E-infrastructures are geographically distributed computing resources and data storage facilities linked by high-performance networks. They allow scientists to share information securely, analyse data efficiently and collaborate with colleagues worldwide. They are an essential part of modern scientific research and a driver for economic growth. By 2020, e-infrastructures need to be deployable on demand to provide a foundation for the digital European Research Area (ERA).

The European Grid Infrastructure (EGI) is a collaborative production infrastructure built as a federation of national resource providers to support multidisciplinary science across Europe and around the world. The EGI ecosystem of national and European funding agencies, research communities, technology providers, technology integrators, resource providers, operations centres, resource centres, coordinating bodies and other functions currently support over 20,000 researchers and their intensive data analysis.

The EGI ecosystem provides a common foundation through which the digital ERA can be brought online, building upon the strengths that have been developed over the last decade.

Introduction



During the last year, the strategic development of EGI.eu and its role in the ERA have been discussed in depth in order to keep up with the rapidly evolving environment.

In this context, the most important driver of change is the European Strategy Forum on Research Infrastructures (ESFRI) projects and their effort to broaden the scope of European science. The ESFRIs present a number of new challenges and opportunities for EGI as an e-infrastructure to support collaboration and all forms of computing.

One characteristic of many of the ESFRI research infrastructures is that they are distributed and data intensive. Their requirements typically include support to connect different types of data sources, experiments and computing resources, all made available to individual users through flexible, but strongly secured services.

EGI.eu is already responding to this demand, by prototyping a European-wide cloud computing capability for researchers by federating the private clouds being established in EGI's resource centres across Europe for their use.

In the forthcoming years, EGI.eu and its infrastructure will continue to evolve and to prove its capability to implement the latest technologies and meet the ever changing demands of leading European research communities.

Per Öster
Chair, EGI.eu Executive Board and EGI Council



Welcome to this year's annual report of the European Grid Infrastructure and its coordination body EGI.eu.

During our second year of operation, many of EGI.eu's activities on behalf of the community have successfully become established as routine operations. This has allowed us to start focusing beyond our current activity on new horizons: our engagement with new research communities to understand their needs for distributed computing infrastructures and the roles EGI can play, and to plan as a community our ambitions and goals for 2020.

Over the last year EGI.eu has also started to engage in collaborations with other projects representing aspects of the EGI community. These range from the environmental research infrastructures and biodiversity communities in the ENVRI and BioVel projects, to economic assessments of e-infrastructure in the e-Fiscal project. The past few months have seen negotiations conclude on projects relating to the digital cultural humanities, workflow systems, cloud computing and federated service management that will all bring in external expertise, new research communities and consolidate the current activity taking place in EGI.

The increasing strength of our interactions with the community and the continuing support that they offer is clear evidence of the role EGI.eu can play in providing added European value in the years to come.

Steven Newhouse
Director, EGI.eu

EGI is a federation of resource providers set-up to deliver sustainable, integrated and secure computing services to European researchers and their international partners.



EGI.EU PARTICIPANTS†

Armenia: NGI-AM, RAMES SCS
 Belgium: NGI-BE, BELNET
 Bosnia and Herzegovina:
 NGI-BA, ETFBL
 Bulgaria: NGI-BG, IPP-BAS
 CERN
 Croatia: NGI-HR, SRCE
 Cyprus: NGI-CY, CyGrid
 Czech Republic:
 NGI-CZ, CESNET
 Denmark: NGI-DK, DCSC

Estonia: NGI-EE, EENet
 Finland: NGI-FI, CSC
 France: NGI-FR, CNRS
 Germany: NGI-DE, Gauß-Allianz
 Greece: NGI-GR, GRNET
 Hungary: NGI-HU, NIIF
 Ireland: NGI-IE, Grid-Ireland
 Israel: NGI-IL, IUCC
 Italy: NGI-IT, INFN
 Latvia: NGI-LV, LUMII

Lithuania: NGI-LT, VU
 Luxembourg: NGI-LU, RESTENA
 FYR Macedonia: NGI-MK, MARGI
 Moldova: NGI-MD, RENAM
 Montenegro: NGI-ME, UoM
 The Netherlands: NGI-NL, NCF
 Norway: NGI-NDGF,
 UNINETT Sigma AS
 Poland: NGI-PL,
 CYFRONET AGH
 Portugal: NGI-PT, UMIC

Romania: NGI-RO, ICI
 Russia: NGI-RU, e-ARENA
 Serbia: NGI-RS, IPB
 Slovakia: NGI-SK, SlovakGrid
 Slovenia: NGI-SI, ARNES
 Spain: NGI-ES, CSIC
 Sweden: NGI-SE, SNIC
 Switzerland: NGI-CH, SWITCH
 Turkey: NGI-TR, ULAKBIM
 United Kingdom: NGI-UK, JISC

† in alphabetical order by [country, if applicable]: [NGI code used in this document, see page 32], [short name of the lead organisation]

What does EGI do?

The European Grid Infrastructure (EGI) was established in 2010 as a European-wide federation of national computing and storage resources, to support cutting-edge research, innovation and knowledge transfer in Europe.

Building on over a decade of investment by national governments and the European Commission, EGI supports 20,000 researchers across many fields of research with a wide range of technical and infrastructure services. EGI's services are distributed across Europe and beyond over more than 350 resource centres, supporting in excess of 1.2 computing million jobs per day, as well as data storage, transfer and open access.

EGI is coordinated by EGI.eu, a not-for-profit foundation supported by the EGI-InSPIRE project (until 2014) and governed by the national stakeholders and early-adopting international research communities.

Vision

To support the digital European Research Area through a pan-European research infrastructure, based on an open federation of reliable services, which provide uniform access to national computing, storage and data resources.

Mission

To connect researchers from all disciplines with the reliable and innovative ICT services they need to undertake their collaborative world-spanning research.

Core Values

Leadership: EGI is a leading pan-European infrastructure, integrating worldwide computing, storage and data resources to support an ecosystem built on innovation and knowledge transfer.

Openness: EGI operates collaboratively with a transparent governance structure that integrates the views and the requirements of all stakeholders, from research communities to resource providers.

Reliability: EGI provides a reliable infrastructure that research communities can depend on in order to collaborate with their peers and deliver innovation.

Innovation: EGI will continue to meet the needs of research communities operating at unparalleled geographic and technical scale by partnering to bring new technologies into production.



EGI IN NUMBERS

- EGI Council participants: 38
- EGI-InSPIRE project partners: 50 (from 56 countries)
- Resource providers: 42
- 20,883 users in 226 virtual organisations
- CPU cores (EGI-InSPIRE partners): 270,800
- Computing capacity (EGI-InSPIRE partners): 2.96 million HEP-SPEC 06
- Usage: 41 million jobs/month
- Reliability: 95.42%
- Availability: 94.50%

Availability of a service or a site is defined as: the time the service/site was up and running, divided by a given period (during which the service was tested), expressed as a percentage.


Reliability of a service or a site is defined as: the actual time the service/site was up and running, divided by its expected run time, in a given period (during which the service was tested), expressed as a percentage.

Registered grid users by research discipline:



Total number of users = 20,883 (April 2012)

Usage of NGI grid computing resources:



EGI's European impact

EGI's vision, mission and core values ensure that it continues to build on its European added value, supporting world class research and innovation in the ERA. Some examples of EGI's impact on science and society are presented here.

SCIENCE AND SOCIETY

Case Studies

MODELLING EARTHQUAKES IN THESSALONIKI

Predicting earthquakes is impossible, but grid computing helps to anticipate their effects.

Earthquakes are amongst the world's most destructive natural disasters. The Fukushima-Daiichi nuclear disaster in Japan, for example, was caused by an earthquake-triggered tsunami. It's impossible to predict when or where earthquakes will strike. The only way to avoid catastrophes is good planning and damage control.

Earthquakes are caused by seismic waves that travel through solid rock like ripples on a pond, causing shakes and tremors as they go. Every region has its own geological setting (different rocks, with a different structure) that influences how seismic waves propagate within it. So by studying how seismic waves travel through a particular geological setting, it's possible to predict how a site will respond to an earthquake. This can later be used to understand how constructions could be affected when disaster strikes.

But how do you study seismic wave propagation? One way to do it is to create a virtual earthquake by solving millions of time-sequenced partial differential equations on a given volume of bedrock. This is called the Finite-Difference method. And to be useful it has to take

into account geological and mechanical properties, and the way they change within a region.

Despite the sophisticated name, a partial differential equation is not complex and can be solved with pen and paper. But to solve millions of these equations you need more than a laptop or two – you need grid computing.

Using the grid computing resources made available by the Greek National Grid Initiative, scientists were able to create the theoretical ground motion for the Thessaloniki region. Without these resources, it would have taken months to complete the task.

The data was then used to calculate the peak ground acceleration and the peak ground velocity throughout the region. These two parameters are extremely useful to assess earthquake hazard and potential structural damage in buildings.

<http://go.egi.eu/thess>

HUNTING FOR NEW TYPES OF VIRUSES

How grid computing helps scientists to discover new diseases.

Respiratory infections are the main reason why children under five end up in hospital. However, in up to 40% of the cases it's not possible to define the exact cause of the disease and this means that there are viruses still unknown to science.

Identifying as many viruses as possible improves the chances of correct diagnosis and helps to determine the best treatment for patients. Knowing which virus is responsible for which disease is also very important to detect potential epidemics or to assess the seriousness of viral infections.

The question is: how do we find new species of virus?

Scientists at the Academic Medical Centre of the University of Amsterdam have been working on VIDISCA-454 – a method to spot new viruses from previously unidentified genetic sequences.

The result of a VIDISCA-454 analysis of a patient sample is a haystack of information that includes – somewhere – the genetic sequence of the unknown virus. Finding this genetic needle is difficult. One way of concentrating the search is to compare the mystery sequences to known viruses catalogued in massive reference databases, such as GenBank.

Customised workflows, enabled by grid computing services, define the sequence of computational steps required to perform an analysis. With tools available in online platforms, scientists can analyse a VIDISCA-454 experiment quickly. For example, a test with 1444 samples produced 4,783,684 genetic sequences and was analysed in 14 hours. You would need 17 days to run it on a local server.

<http://go.egi.eu/virus>


HOW FAST COULD A T-REX RUN?

How grid computing is helping paleontologists to understand how dinosaurs moved and what roles they played in their ancient world.

With its sharp teeth and massive jaws, the T-rex is the stuff of nightmares. Scientists are convinced the T-rex was a carnivorous predator but huge teeth don't tell the whole story. Did it catch prey in sprints like a cheetah? Or was it a stalk-and-ambush hunter like the jaguar? What was its place in the Cretaceous ecosystem?

Since we can't see a real T-rex in action (it disappeared along with the other dinosaurs 65 million years ago), palaeontologists need to look elsewhere to understand its role as a predator. Top running speed offers good clues to solving this mystery – but how do you measure the maximum speed of an extinct animal?

The solution is to create a computer simulation of the animal's skeleton and muscles. Such simulations are only useful if they are detailed. Accurate models are computationally demanding and impractical to complete using a single computer. The grid services provided by the UK's NW-Grid solved this problem and helped the researchers to use about 170,000 hours of computing time to complete the project in just a few months.



Is this a new species of virus? Scientists can use the grid to find out.

The results, published in the *Proceedings of the Royal Society B*, show that the T-rex could only speed up to 29 km/h. This means that Usain Bolt – the world’s fastest man – could probably outrun it with his 9.58s 100-metre record (ca. 37.5 km/h).

It’s with grid-enabled models such as this that we know that the T-rex was certainly not the sprinting cheetah of its time and hunted perhaps more like a stalking jaguar.

<http://go.egi.eu/t-rex>

PROTECTING PORTUGAL’S AVEIRO LAGOON

How grid computing allows for a better management of coastal resources.

The Aveiro Lagoon in Portugal is a national treasure. With a length of about 45km and separated from the Atlantic Ocean by a sandy dune barrier, this shallow lagoon is one of Europe’s last coastal marshes and a haven for many bird species. The lagoon is also an important source of revenue in the region, fuelling tourism, aquaculture industries, artisan fishing and the collection of ‘fleur de sel’, a prized variety of salt.

In the past years the lagoon has been threatened by a decrease in water quality. Thanks to the lagoon’s economic, ecologic and cultural importance, there is a strong push to preserve its ecosystem. The key to long-


term sustainability is efficient management and to achieve that, decision-makers need to have a solid understanding of this environment.

Scientists from the National Laboratory for Civil Engineering, Portugal applied a three-dimensional computational model called ECO-SELFE to the Aveiro Lagoon scenario. ECO-SELFE simulates physical variables (e.g. currents, water temperature or salinity), biochemical processes (e.g. carbon, nitrogen cycles) and ecological relationships (e.g. plankton mortality or availability of prey).

The output is a complex but detailed simulation of the Aveiro Lagoon ecosystem. ECO-SELFE is very demanding in terms of computing power and the team turned to the grid computing resources provided by INGRID, the Portuguese National Grid Initiative, for help. Grid computing allowed them to improve computational efficiency significantly, saving a lot of time: the analysis was complete in about a month.

The results, published in the *Journal of Coastal Research*, show that phytoplankton concentration is the ‘canary in the mine’ for the Aveiro Lagoon ecosystem – if there is a sudden or drastic change in the phytoplankton, action needs to be taken as soon as possible.

<http://go.egi.eu/eco-selfe>



Stories from the grid

As part of the new EGI outreach strategy, the communications team commissioned 'Een van de Jongens', a film production company, to produce a series of videos about grid computing applications for science. The result of this partnership is the first three episodes of the 'Stories from the grid' series.

Each five-minute episode focuses on a specific application of grid computing in a different research field. The videos were filmed as documentaries and feature the people actively involved in the research explaining the goals and outcomes of the work.

The videos are freely available to the public on YouTube and on the EGI website.



EPISODE 1 – THE CONE SNAIL

In its natural habitat, the cone snail deploys powerful toxins to paralyse the fish it eats. In the lab, researchers are studying how these toxins can be used to block pain signals to the brain. This work could lead to a new generation of anaesthetics for hospital patients and could also alleviate the muscle spasms caused by the condition dystonia.

Alexandre Bonvin (above) and Henry Hocking from the University of Utrecht explain how grid services lend the extra computing power they need for efficient analysis and interpretation of the data. Without grid computing, the team would need months to perform all the necessary calculations and the whole research project would be unfeasible.

Released February 2012

<http://go.egi.eu/conco>



EPISODE 2 – THE EPIGONION

The epigonion was the guitar of ancient Greece but since none survived the passing of time, it hadn't been heard for centuries. Until now.

Using a technique called physical modelling, Domenico Vincinanza (above) recreated the sound of the instrument's 48 strings as digital files. With the help of grid computing, it took him just a few hours. In a single core computer he would need a month. The epigonion's sounds can now be downloaded and played by any musician using a simple keyboard.

Released May 2012

<http://go.egi.eu/epigonion>



EPISODE 3 – THE TOP QUARK

The Large Hadron Collider (LHC) is the world's largest and most complex experiment, at the cutting edge of high energy physics. Particle physicists use the LHC to study variations from the Standard Model and discover potential new laws of physics. The particle known as the top quark is a window to this weird and wonderful world.

The LHC produces enormous amounts of data, enough to fill piles of DVDs. Marcel Vreeswijk and Hurng-Chun Lee from NIKHEF explain how customised grid computing workflows are key to filtering and sieving the dataset down to a manageable size. Without these tools, it would be impossible to pick out the collision event that could hold the clues to top quark behaviour.

Released June 2012

<http://go.egi.eu/topquark>

EUROPEAN INNOVATION

The next generation of European research innovation will increasingly rely on both geographical and intellectual collaboration across multiple disciplines. Recognising this trend and opportunity, the European Commission has unveiled Europe 2020 – a strategy to develop a smart, sustainable and inclusive economy based on innovation and knowledge transfer.

One of the driving forces behind this strategy is the digital ERA – a transnational effort combining research centres with ongoing programmes and projects aimed at building multidisciplinary collaborations and enabling rapid knowledge transfer across borders.

The full implementation of the digital ERA depends heavily on the development of e-infrastructures that will enable the accomplishment of the ‘fifth freedom’ – free circulation of researchers, knowledge and technology across Europe. EGI will be a key enabler of this vision by linking thousands of researchers across Europe to the resources they need, in collaboration with other pan-European e-infrastructures such as GÉANT (for networking) and PRACE (for supercomputing).

In practice, EGI’s contribution to the Europe 2020 strategy will focus on the Digital Agenda for Europe and Innovation Union flagship initiatives, as well as providing an underlying added value that will bring cross-cutting benefits to European research and society as a whole.

EGI and the Digital Agenda for Europe

In support of the Digital Agenda for Europe, EGI will:

- Provide a single uniform market for accessing distributed computing resources and connected data in Europe through EGI’s federation of national resource providers.
- Promote competitiveness and interoperability for e-infrastructures through open standards.
- Reduce inefficient research spending and stimulate innovation across Europe by maximising the use of national resources and knowledge within a common federated infrastructure and community.
- Offer large scale ICT facilities that enable the exploration of new computing and data processing models that address scientific grand challenges facing society.

EGI and the Innovation Union

In the context of the Innovation Union, EGI can:

- Position itself as a key enabler of the digital ERA that will allow for the free circulation of researchers, knowledge and technology.
- Promote excellence in education and skills development by simplifying multi-disciplinary cooperation.
- Bridge geographical boundaries beyond Europe thanks to the many collaborations and integration with worldwide e-infrastructures.

Added value to Europe

In the wider context of Europe 2020 (funded in future through the Horizon 2020 programme), EGI offers a range of added value to Europe:

- EGI ensures that researchers have access to uniform and reliable computing resources, enabling faster scientific results and avenues of multi-disciplinary research otherwise not possible.
- EGI allows researchers to focus on their research rather than managing their e-infrastructures needs.
- EGI's monitoring and operational services guarantee an integrated, reliable and uniform service provided across organisational and national boundaries.
- EGI's transnational integration provides effective use of resources and ensures the most effective return on Europe's e-infrastructure investments.
- EGI enables knowledge transfer and sharing of solutions through community events, an application database and a training marketplace.
- EGI promotes open science through the availability, accessibility and reuse of scientific data and results, the provision of web-based tools that facilitate scientific collaboration and by ensuring public access to research.

EGI policy papers

- EGI.eu Position Paper on Common Strategic Framework
<http://go.egi.eu/664>
- EGI Position Paper on ERA Framework
<http://go.egi.eu/891>

Supporting the EGI ecosystem

Building on the experiences acquired during the European DataGrid (EDG) and the Enabling Grids for E-Science (EGEE) projects, the EGI ecosystem is now emerging as a thriving community of independent technology providers, platform integrators and operators, resource infrastructure providers, research communities, as well as national (NGIs) and European (EGI.eu) coordination bodies.

EGI.eu's mission is to nurture this environment and develop it into a sustainable ecosystem, able to maintain a consistent and reliable level of ICT services for European researchers and their international colleagues. To ensure this, the different components of the ecosystem need to be developed and supported individually, with effective processes through governance structures and partnerships that enable them to scale as required.

OPERATIONS MANAGEMENT

The EGI.eu distributed Operations Team coordinates and performs the activities required to deliver services at an agreed level. Operations is also responsible for the ongoing management of technology deployment and technical support services. This includes the monitoring and accounting, grid oversight, service level management, staged rollout of new software releases and integration of new software platforms.

Engagement with new infrastructure providers continues to be a priority and during 2011–2012 two Memoranda of Understanding (MoUs) were signed with resource infrastructure providers from South Africa and Ukraine. The tools and services required to manage the infrastructure continue to be developed. These include the Operations Portal, the infrastructure technology database (GOcdb), the accounting and metrics portals, the EGI Helpdesk, the Service Availability Monitor, the Applications Database and the Training Marketplace.

SUPPORT AND COMMUNITY COORDINATION

The EGI.eu User Community Support Team (UCST) supports users of the infrastructure through close cooperation with NGIs and VRCs. Such support focuses on creating applications, tools, knowledge and events that can be adopted and tailored by the community to their own particular needs.

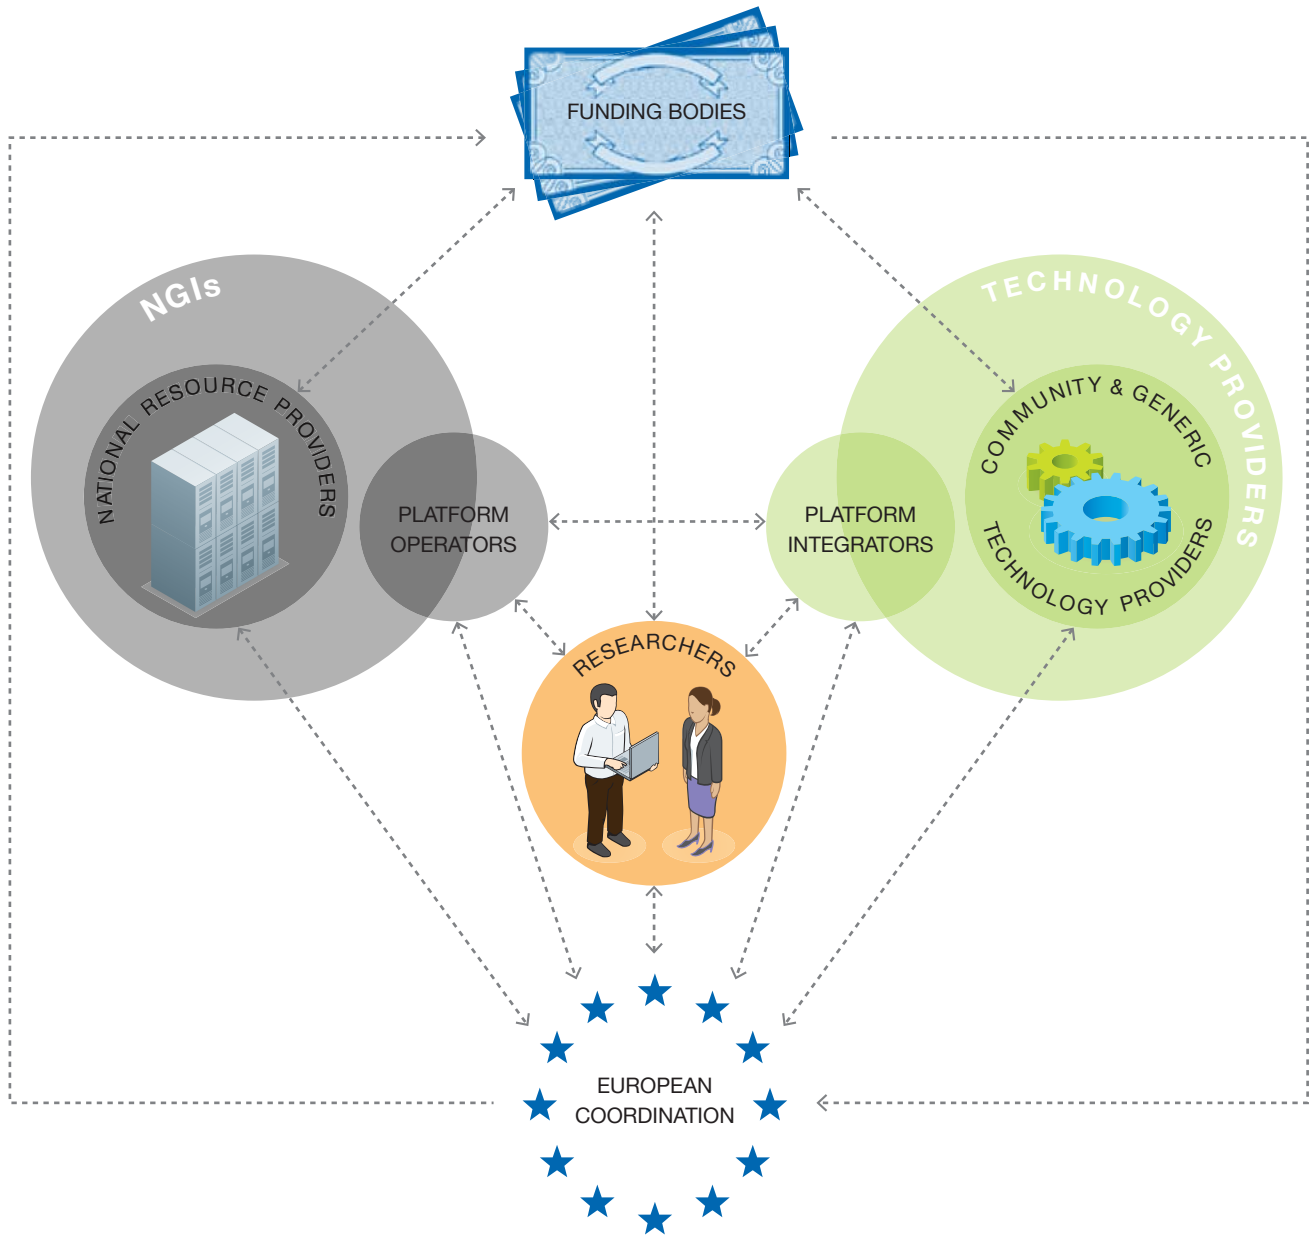
A second, but equally important, focus is the bi-directional nature of the communication between UCST, mainly through the User Community Board (UCB), and the different research communities: apart from acting as a forum for the research communities involved in EGI, the UCB also listens to their needs, feeding important requirement information back to the operations and technical teams looking after the infrastructure.

The EGI community continued to grow in 2011–2012, with MoUs signed with three important Virtual Research Communities:

- The Life-Science Grid Community (27 May 2011)
- The Hydro-Meteorology Research Community (22 August 2011)
- The Worldwide LHC Computing Grid (24 January 2012)

Another important milestone was the signing of a Letter of Intent with the CLARIN (Common Language Resources and Technology Infrastructure) and DARIAH (Digital Research Infrastructure for the Arts and Humanities) ESFRI projects. This marks a significant step towards the goal of expanding the usage of grid infrastructures to the humanities field, well beyond the early-adopter research communities.


Overview of the EGI ecosystem:



TECHNOLOGY PROVISION

All software deployed on EGI is produced through independent platform integrators and technology providers. EGI.eu helps manage the outsourcing of technology developments through the Technology Coordination Board (TCB), according to the requirements gathered by the UCB and the Operations Management Board (OMB), negotiates with potential technology providers and assesses the quality of the new software.

Over the second year, the Technology coordination of EGI.eu consolidated and improved the processes that were designed and implemented during the first year. The release of the EMI-1 middleware package in early May 2011 was the biggest challenge for the Software Provisioning infrastructure and was well integrated with the Staged Rollout team, who managed the deployment of the software into production.



EGI FEDERATED CLOUD TASK FORCE

The EGI Federated Cloud Task Force was established in July 2011 to work with the community to develop a 'blueprint' for EGI resource centres who wish to securely federate and share their local virtualised environments externally with collaborators as part of the production infrastructure. Other goals include:

- Define and prototype solutions for monitoring, accounting and advertising through the information services virtualised resources
- Investigate and analyse requirements from early adopting research communities
- Provide feedback to relevant technology providers (both within and external to the TCB)
- Identify issues that need to be addressed by other areas of EGI (e.g. policy, operations, support and dissemination)

While the work of the Task Force is still in progress, the group has defined a total of eight capabilities required of a future EGI Federated Cloud. These are: Virtual Machine Management, Storage/Data management, Information Discovery, Accounting, Monitoring, Notification, Federated Authentication and Authorisation Infrastructure and Virtual Machine Image sharing.

The state of the art of the Federated Cloud infrastructure was demonstrated at the EGI Community Forum 2012 in Munich.

COMMUNITY ENGAGEMENT

Teams in EGI.eu work with staff in the NGIs through the NGI International Liaisons (NILs), and with the broader EGI community through the Virtual Team (VT) framework, to develop EGI's activities and to engage with research communities new to EGI. The VT framework is a new mechanism to allow interaction, knowledge transfer and cooperation within the community on non-operational activities, covering areas such as communications, policy and strategy, outreach and support, and events.

Strategy, policy and collaborations

The EGI.eu Strategy and Policy Team's (SPT) work focuses on developing strategies and policies, within and external to the EGI ecosystem, relating to governance, standardisation and integration with other infrastructures, strategic response and alignment to EU strategic development.

The main objectives during the year were to analyse global and European themes and trends, to produce reports to inform the EGI community about strategic-level developments, to organise meetings and workshops on strategic themes key to EGI and to nurture and develop collaborations within and outside the EGI ecosystem. The EGI Compendium – another activity that started this year – will track the annual evolution of the EGI community and its NGIs.

WHO IS AN NGI INTERNATIONAL LIAISON?

During 2011, EGI introduced a new role in the community – the NGI International Liaison (NIL). Working together with teams in EGI.eu and other NGIs, the NILs link the strategic non-operational activities taking place in their own NGIs (e.g. outreach, marketing, communication, training, new community engagement) with the rest of the community.

With the NIL acting as a spearhead, each country will be able to demonstrate their added value at a European level and to share their skills with the whole community.




Events

TECHNICAL FORUM 2011

The EGI Technical Forum was held from 19–23 September 2011 in Lyon, France, to drive forward progress toward the adoption of a federated virtualised infrastructure for European researchers. The Technical Forum was organised by EGI.eu in conjunction with the French NGI and the CC-IN2P3 computing centre.


One of the added values of the EGI Forums is the opportunity to co-locate related events. In this instance these included the 9th e-infrastructure Concertation Meeting, OGF31, Grid2011 in addition to the first French Grid Day and the first GlobusEUROPE event which was organised by the IGE project. Other EU-funded projects such as SIENA also organised workshops under this umbrella.

In total, 655 participants registered for the forum, making it the best attended EGI event so far. The sessions included 132 contributions. The Technical Forum website (<http://tf2011.egi.eu>) received almost 4,000 unique visitors, with about 35,000 page views.



Technical Forum 2011:

- poster to promote the Forum
- the Forum venue—the Cité Centre de Congrès in Lyon
- a session discussing the model for the EGI Roadshow





Community Forum 2012:



1. exhibition booths
2. opening keynote by Kostas Glinos, Head of the 'GÉANT & e-Infrastructure' unit at the EC
5. the Forum venue – the EGI booth can be seen in the top right corner, with two EGI demo booths in the foreground.
7. poster to promote the Forum
8. the keynote and plenary sessions were well-attended

COMMUNITY FORUM 2012

The EGI Community Forum was held from 26–30 March 2012 in Munich, Germany, in conjunction with the German NGI and the Munich Network Management team. The event showcased the role that EGI plays in enabling innovation across the European Research Area, in particular by highlighting the services, technologies and tools available to, and also provided by, scientific communities to support their research.

The Community Forum was held in conjunction with the second EMI Technical Conference and co-located with the second Annual European Globus Community Forum (26 March). In total, 421 people registered for the event that featured 171 contributions. The Community Forum website (<http://cf2012.egi.eu>) was read by about 3,200 unique visitors, with about 20,000 page views.

Social media channels were very active during the conference, enhancing discussions and community-building. The #cf2012 tag was used in 403 Twitter posts by 60 different users, while 221 photographs were uploaded to Flickr with the tag egicf12. The two (iPhone and Android) free smartphone applications developed for the event were downloaded by 154 users.



Publications

During 2011–2012, the EGI.eu communications team produced several publications designed to convey the goals and achievements of the EGI community to wider audiences. These include:

- Four issues of the *Inspired* newsletter
- 80 news items, on topics as varied as events, announcements, milestones and agreements, among others
- The EGI-Public Service Review brochure and cover DVD, explaining the goals of EGI to an audience of 140,000 decision makers
- The *Earth Science* brochure, with a selection of case studies focusing on earth sciences, for the European Geophysics Union General Assembly attended by 11,000 delegates. In addition, a *Life Science* brochure was produced with relevant case studies, to respond to the need of a growing community.
- The EGI brochure, aimed at introducing EGI to a general audience
- The *EGI Community Forum 2012 Book of Abstracts* (ISBN 978 90 816927 00), with an overview of the talks presented at the event
- Two Gender Action Plan postcards, to raise awareness of gender issues in the ICT community
- The Training Marketplace bookmark, to announce the launch of the knowledge transfer scheme

EGI and the work done on the infrastructure featured in a wide range of articles and press cuttings including *Public Service Review*, *Pan-European Networks*, *The Parliament*, *Symmetry*, *Discovery News*, *Wired* and a mention in the UK's *Daily Mail* social media channel.

Website and social media

EGI relaunched the website in March 2012 with improved graphics and navigation and dedicated channels for different audiences, such as users, researchers, policy makers and the general public. The site also showcases the new Facebook pages and Twitter feeds, which have been used to launch competitions, teaser campaigns and videos. These complement the EGI blog and newsfeed as forums to allow the community to know what the community is working on, and participate in the conversation. The new social media strategy offers an improved presence with a dedicated voice on Facebook and Twitter to three EGI.eu teams: policy and strategy, user community support and a combined channel for technology and operations.

These teams now have tools to engage directly with their community, complementing EGI information provided through the generic channel. Bringing together the news, information and discussions in each area means that everyone can get involved and understand where EGI is going. These new channels will allow the teams to report from meetings, events and discussions as they happen, so everyone can get engaged in the process and have their say about the issues that matter to them.

- Website: www.egi.eu
- Social media channels: <http://go.egi.eu/smc>



Above: the new EGI website, launched during the Community Forum 2012.


Opposite page (clockwise from top left): the EGI brochure; the 'EGI Community Forum 2012 Book of Abstracts'; 'Life Science' brochure; EGI/Public Service Review joint publication; 'Inspired' newsletter; 'Earth Science' brochure.

TECHNICAL OUTREACH TO NEW COMMUNITIES

Converting a research community that is not using EGI into an active user community within EGI requires technical effort and planning to accommodate their individual needs. Each new group is different in its nature and requirements. This means that for each new community it is important to identify which resources they need, to consider the integration of new resources into EGI, to port applications to an EGI platform, to deploy services to meet new requirements and provide training.

Recognising the importance of new communities to the ecosystem, EGI.eu has created a new team dedicated to technical outreach to new communities. The team works closely with NGI International Liaisons to ensure that a coordinated, systematic and strategic approach is taken to technical outreach, including gathering requirements.

Another important aspect of this community-building effort is to ensure that new users make the most of the expertise and the tools accumulated and developed by EGI and its predecessors. To this effect, the technical outreach to new communities also promotes knowledge transfer from within and between the early-adopter heavy user communities; a training marketplace; an application database; a list of NGI assets and activities and engagement in external collaborations.



EGI's strategy for the future

EGI today is the result of pioneering work that has built a pan-European production infrastructure through the federation of national resource providers, to support multi-disciplinary science across Europe and around the world. With thousands of users across many fields of science, EGI is already making a positive impact on European research and innovation.

But EGI is not a static resource. The community is already looking towards the future, to identify opportunities, expand its service offering and anticipate challenges.

EGI is committed to the European Commission's goals outlined in the Europe 2020 vision. With this in mind, EGI's strategy for the future is to develop its activities in order to be a key enabling foundation of the digital ERA, supported by continued investment from national and European funding bodies.

In practice, this means developing EGI's strengths in three key areas:

- community-building and coordination
- operational infrastructure
- virtual research environments

COMMUNITY-BUILDING AND COORDINATION

Community-building and coordination experience help to guarantee integration and deliver uniform and open access to computing resources.

EGI has developed from a domain-specific activity focused on a handful of resource centres to a multi-disciplinary infrastructure with usage across many different disciplines. Individual researchers can expect uniform access to reliable services anywhere within EGI's federated network of national resource providers, regardless of location. The value of such widespread integration has been amply demonstrated (e.g. during the data collection runs of the Large Hadron Collider in 2010–2012 where scientists analysed data using EGI services).

EGI is a truly pan-European infrastructure. This means that EGI federates resources from dozens of European countries, each with its own strengths, structures, priorities and procedures. The quality of service offered to scientists today is the result of a decade of cooperation and consensus-building around many non-technical issues, including privacy, governance and strategy.

Therefore, EGI's governing body – the EGI Council – is ideally placed to build on these achievements and lead the network of NGIs and EIROs to the next level. With the coordination provided by EGI.eu, EGI will continue to develop as an infrastructure and, as importantly, as a community through communication, marketing, outreach, support and public events.

OPERATIONAL INFRASTRUCTURE


Developing the tools and services to support existing communities and to offer virtualised resources through a federated cloud.

EGI's success in providing reliable services to research communities depends on a sophisticated, infrastructure-wide network of services. These operational services allow for the monitoring of individual resource centres, the collection of accounting records nationally and across the whole infrastructure, the resolution of technical issues across different organisational structures, and the compilation of availability and reliability statistics.

EGI's mature monitoring, accounting and service offer was not born in a day – it is the fruit of the knowledge and expertise built up and shared by individual resource centres within their NGIs and brought together within the community. EGI's decade-long experience in managing a federated infrastructure is a key asset and a service that can be offered to any research community seeking to build a European-wide operational infrastructure.

This experience will now be used to support services being run in other data centres for other research communities, and to create a federated cloud infrastructure that will complement the EGI service offering with virtualised resources.

Virtualisation technologies are already being adopted by many resource centres through private clouds (accessible only to internal users) that improve management and give greater flexibility to local users. Up to now, no clear model has emerged on how to maximise private cloud infrastructures through federation and integration so as to benefit non-local users. EGI provides a perfect environment to tackle this problem by developing a federated cloud infrastructure built around open standards.



VIRTUAL RESEARCH ENVIRONMENTS

Customising and deploying Virtual Research Environments to support world-class science.

Virtual Research Environments (VRE) are made up of the applications, services and tools that are deployed across Europe between the researcher and the e-infrastructure. Initially, VREs were developed as command-line interfaces, simple to use with the right training.

Over the years, the EGI community of users expanded to include researchers with more diverse backgrounds and interests (not always technically-oriented). The diversification of EGI led to the development of higher-level generic tools and domain-specific VREs to simplify the data analysis process. As the democratisation of the infrastructure continues, research communities using EGI in the future will have divergent needs and a requirement to deploy the VRE that they need, where they need it. The only scalable solution is to allow and support customisation.

Science gateways (or portals) are VREs designed to reduce the technical barriers to accessing remote computing resources. This is an area ripe for innovation and already supported by EC and national investment initiatives, alongside an active open-source community and standards activity around generic portal frameworks.

EGI – through its coordinating body EGI.eu – is committed to coordinate VRE development, collecting requirements and organising domain-specific workshops and sessions at the Technical and Community Forums. The use of the web as a route to accessing the e-infrastructure (ranging from desktops, to tablet applications or mobile phone browsers) provides an unprecedented opportunity to meet and anticipate the demand coming from young researchers to access e-infrastructures with the tools available to them during their daily activities.



European projects

EGI-INSPIRE

EGI-InSPIRE (EGI-Integrated Sustainable Pan-European Infrastructure for Research in Europe) is a four-year project helping to establish a sustainable, reliable e-infrastructure that can support researchers' needs for large-scale data analysis. EGI-InSPIRE is a collaborative effort involving 50 partners in over 40 countries.

The project is co-funded by the European Commission's 7th Framework Programme (contract number: RI-261323) to help lay down the EGI operational and support processes, as well as to build a sustainable e-infrastructure, independent from project cycles.

By the time EGI-InSPIRE finishes in 2014, EGI will be a dependable and more sustainable provider of computing resources for European scientists and researchers.

Work during the second year of EGI-InSPIRE focused on addressing the recommendations coming from EGI-InSPIRE's first review in June 2011. The reviewers' feedback provided a timely opportunity to reflect on the community's strengths and weaknesses and to consider the opportunities and threats EGI might be likely to see in the next decade of grid and cloud computing.

The subsequent discussions about EGI's past, present and – most importantly – future are now summarised in the report 'Seeking New Horizons: EGI's Role in 2020' (<http://go.egi.eu/EGI2020>) (PDF, 1MB). While continuing to deliver and evolve our current production infrastructure, an additional focus during the remainder of the project will be to prepare for the European Commission's Horizon 2020 funding programme, which starts during our final year. Our future focus will be on the community and coordination activities needed to create an open EGI ecosystem, maintaining and developing the Operational Infrastructure and assisting research communities to build and operate the customised Virtual Research Environments they need for their research work.

The activities within the community and coordination area are focused on developing policy, communications and outreach within and by the NGIs, by building on the network of NGI International Liaisons (NILs). Working in conjunction with EGI.eu staff, the NILs bring the resources and skills of their NGI into virtual teams that are able to tackle critical issues within the EGI community. This is being achieved through short-lived projects in areas that will bring in new research communities to EGI, such as developing targeted marketing material and helping to define community-wide policies.

KEY RESULTS

Project Overview

- EGI-InSPIRE presentation <http://go.egi.eu/1145>
- EGI-InSPIRE paper <http://go.egi.eu/201>

Community

- EGI 2011 Technical Forum, Lyon, France
<http://tf2011.egi.eu>
- EGI 2012 Community Forum, Munich, Germany
<http://cf2012.egi.eu>
- EGI 2012 Community Forum Book of Abstracts
<http://go.egi.eu/CF12BoA> (PDF file 4.3MB)
- Work of the Asia Pacific Region <http://go.egi.eu/1067>

Technical Architecture

- EGI Technical Roadmap
<http://documents.egi.eu/document/1094>
- EGI Operations Architecture <http://go.egi.eu/763>

General Reports

- EGI's strategy <http://go.egi.eu/EGI2020> (PDF file 1MB)
- Position paper on the ERA Framework
http://go.egi.eu/EGI_ERA
- Evolving EGI's Business Model <http://go.egi.eu/1040>
- EGI Sustainability Plan <http://go.egi.eu/1147>

EGI-InSPIRE Annual Activity Reports

- External relations <http://go.egi.eu/1069>
- Production infrastructure <http://go.egi.eu/1059>
- Software provisioning and support <http://go.egi.eu/1015>
- Heavy user communities tools and services
<http://go.egi.eu/742>
- Operational tool maintenance and development
<http://go.egi.eu/1063>

E-SCIENCETALK

The e-ScienceTalk project (contract number: RI-260733) started on 1 September 2010 with the mission of bringing the success stories of Europe's e-infrastructure to a wider audience. The project coordinates the dissemination outputs of EGI and other European e-infrastructure projects, ensuring their results and influence are reported in print and online. The outputs of the project are:

- e-ScienceBriefings, key reports for policy-makers
- e-ScienceCity an introduction to e-science for a general audience, including the award winning GridCafé
- *International Science Grid This Week (iSGTW)*, an online weekly newsletter about e-Science
- GridGuide, an atlas of grid computing sites around the world
- GridCast, a blog that goes behind the scenes of e-Science events
- Real Time Monitor, which shows traffic on the grid in near real time

In the second year of e-ScienceTalk, three policy briefings have been published about desktop grids, global connectivity and visualisations in science. The e-ScienceCity website was launched, including new areas about cloud computing and volunteer computing, together with an associated 3D virtual world. Subscribers for *iSGTW* have grown to 8000, with 3000 followers on social media. A dedicated *iSGTW* editor now reports from the Asia Pacific region, and articles have featured interviews with Nobel prize winner Brian P. Schmidt, former UK prime minister Gordon Brown and web pioneer Tim Berners-Lee, and have been picked up by other publications including *Symmetry*, *Discovery News* and *Wired US*, which has 1.4 million Twitter followers. GridGuide has expanded by 25 sites in Europe, the US, Africa and Asia. GridCast has reported from more than a dozen events worldwide and GridCast videos have been viewed nearly 100,000 times in YouTube. The Real Time Monitor now includes jobs from ATLAS, one of the world's largest users of the grid, and shows router traffic on the network layer.

www.e-sciencetalk.org

BIOVEL

The BioVeL consortium (contract number: RI-283493) connects the IT and the biodiversity research fields – two communities with complementary expertise. The IT field will contribute by developing workflow technology and data processing for a virtual laboratory. The experts in the biodiversity field will put such workflows to use in taxonomy, phylogenetics and ecological niche research.

EGL.eu will assist the BioVeL project to develop cloud resources to operate modelling simulations for biodiversity studies that can be used on EGL resources.

www.biovel.eu



E-FISCAL

The main goal of the e-FISCAL project (contract number: RI-283449) is to analyse the costs and cost structures of the European High-Throughput and High-Performance Computing (HTC and HPC) e-infrastructures. These research infrastructures are mainly coordinated and funded by national entities. e-FISCAL will compare them with similar commercially leased or on-demand offerings.

Understanding the overall costs of these European research services is a prerequisite in planning their long-term sustainability, e.g. by new business models for service provision. A quantitative analysis of the cost factors involved will help service providers and user communities to identify areas where the overall cost efficiency of ICT-enabled research can be optimised. The study will also go beyond a simple 'cost per core hour' comparison and it will analyse qualitative differences in service between HTC and HPC e-infrastructures and their closest commercial counterparts.

EGL.eu provides an essential link to a large group of resource centers in the NGIs, which are the targets of this study. As the coordination body of the largest e-infrastructure in Europe, EGL.eu also provides extensive experience in guiding the project's goals.

<http://efiscal.eu>

ENVRI

Cutting edge environmental research increasingly depends on a wide range of data and advanced capabilities to process and analyse them. The ENVRI project, 'Common Operations of Environmental Research Infrastructures' (contract number: RI-283465) is a collaboration in the ESFRI Environment Cluster, with support from ICT experts, to develop common e-science components and services for their facilities. The results will speed up the construction of these infrastructures and will allow scientists to use the data and software from each facility to enable multi-disciplinary science.

EGL.eu's contribution to the ENVRI project will focus on data integration, harmonisation and publication, dissemination and sustainability. The purpose of EGL.eu's participation is to understand how common use cases of the environmental ESFRIs can benefit from existing EGL infrastructural services (e.g. authentication and authorization, accounting, data transfer), so that common solutions can be shared among multiple Research Infrastructures, and to help to communicate the results of this work to both EGL and ENVRI communities.

www.envri.eu

EGL.eu financial statement 2011

The following accounts relate from the establishment of EGL.eu on 1 January 2011 to 31 December 2011. The complete financial accounts and accompanying statement are available online (<http://go.egi.eu/1146>). The accounts were adopted by the EGL Council in June 2012.

BALANCE SHEET

ASSETS		€
Trade and other receivables ¹	1,514,717	
Cash and cash equivalents ²	672,977	
	2,187,694	
LIABILITIES		€
Appropriated reserves ³	994,491	
Current liabilities ⁴	1,193,203	
	2,187,694	

¹ Debtors (€31,484), interest (€62,230), receivables from employees (€776), EGL Technical Forum 2011 (€44,521), project pre-financing (€1,375,706)

² EGL.eu current account (€189,444), EGL.eu savings account (€467,383), deposits (€16,150)

³ Reserve for E-tasks that have not yet been undertaken (EGL.eu: €34,394 and partners –€13,334) and reserves for EGL.eu (€973,431)

⁴ Taxes and social securities (€53,193), project pre-financing (€90,056), advance payments received (€230,855), accounts payable (€47,713), accountant (€5,258), administration (€2,298), other amounts payable (€147), accrued staff pay (€89,675), interests on pre-financing (€27,962) and Global tasks (€646,046).

INCOME

INCOME 2011		€
EGL InSPIRE, 2011 ⁵	1,790,087	
EGL InSPIRE, 2010	343,911	
eScienceTalk	73,823	
EGL.eu participants	965,128	
e-Fiscal	4,359	
ENVRI	1,328	
BioVel	2,179	
Interest	36,776	
TOTAL	3,217,591	

⁵ The income from the EC for the EGL-InSPIRE project has been received in 2011 and includes all adjustments to the 2010 cost statement

EXPENDITURE

EXPENDITURE 2011	€
Staff salaries ⁶	1,621,444
Subsidy for EGI global tasks	655,663
Office costs ⁷	313,855
Direct project costs ⁸	194,102
TOTAL	2,785,064

⁶ EGI.eu has 21 employees at the end of 2011.

⁷ Includes travel costs, fully or partially granted by project.

⁸ Expenditure incurred by EGI.eu on behalf of an EC project that can be reclaimed in full from the EC apart from any VAT paid.

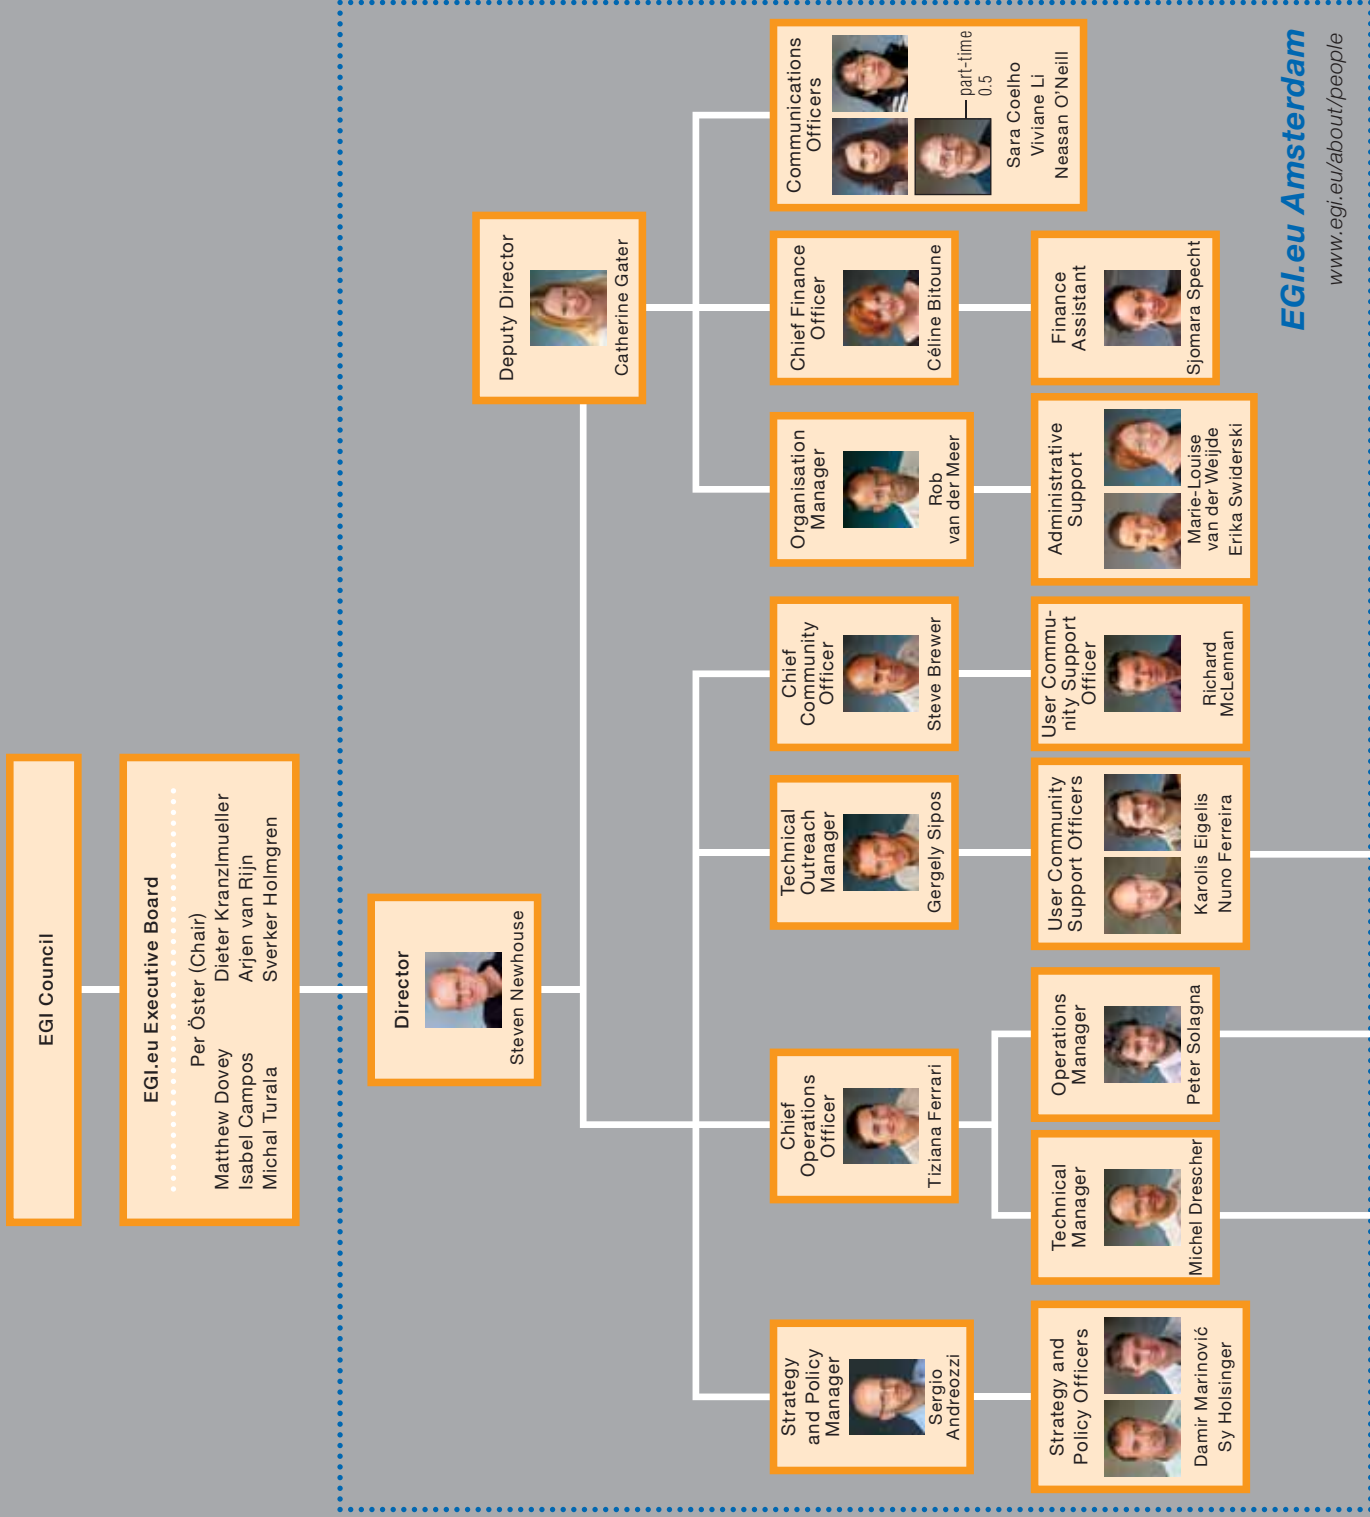
‡ Fee paid for a quarter of the year (25%)

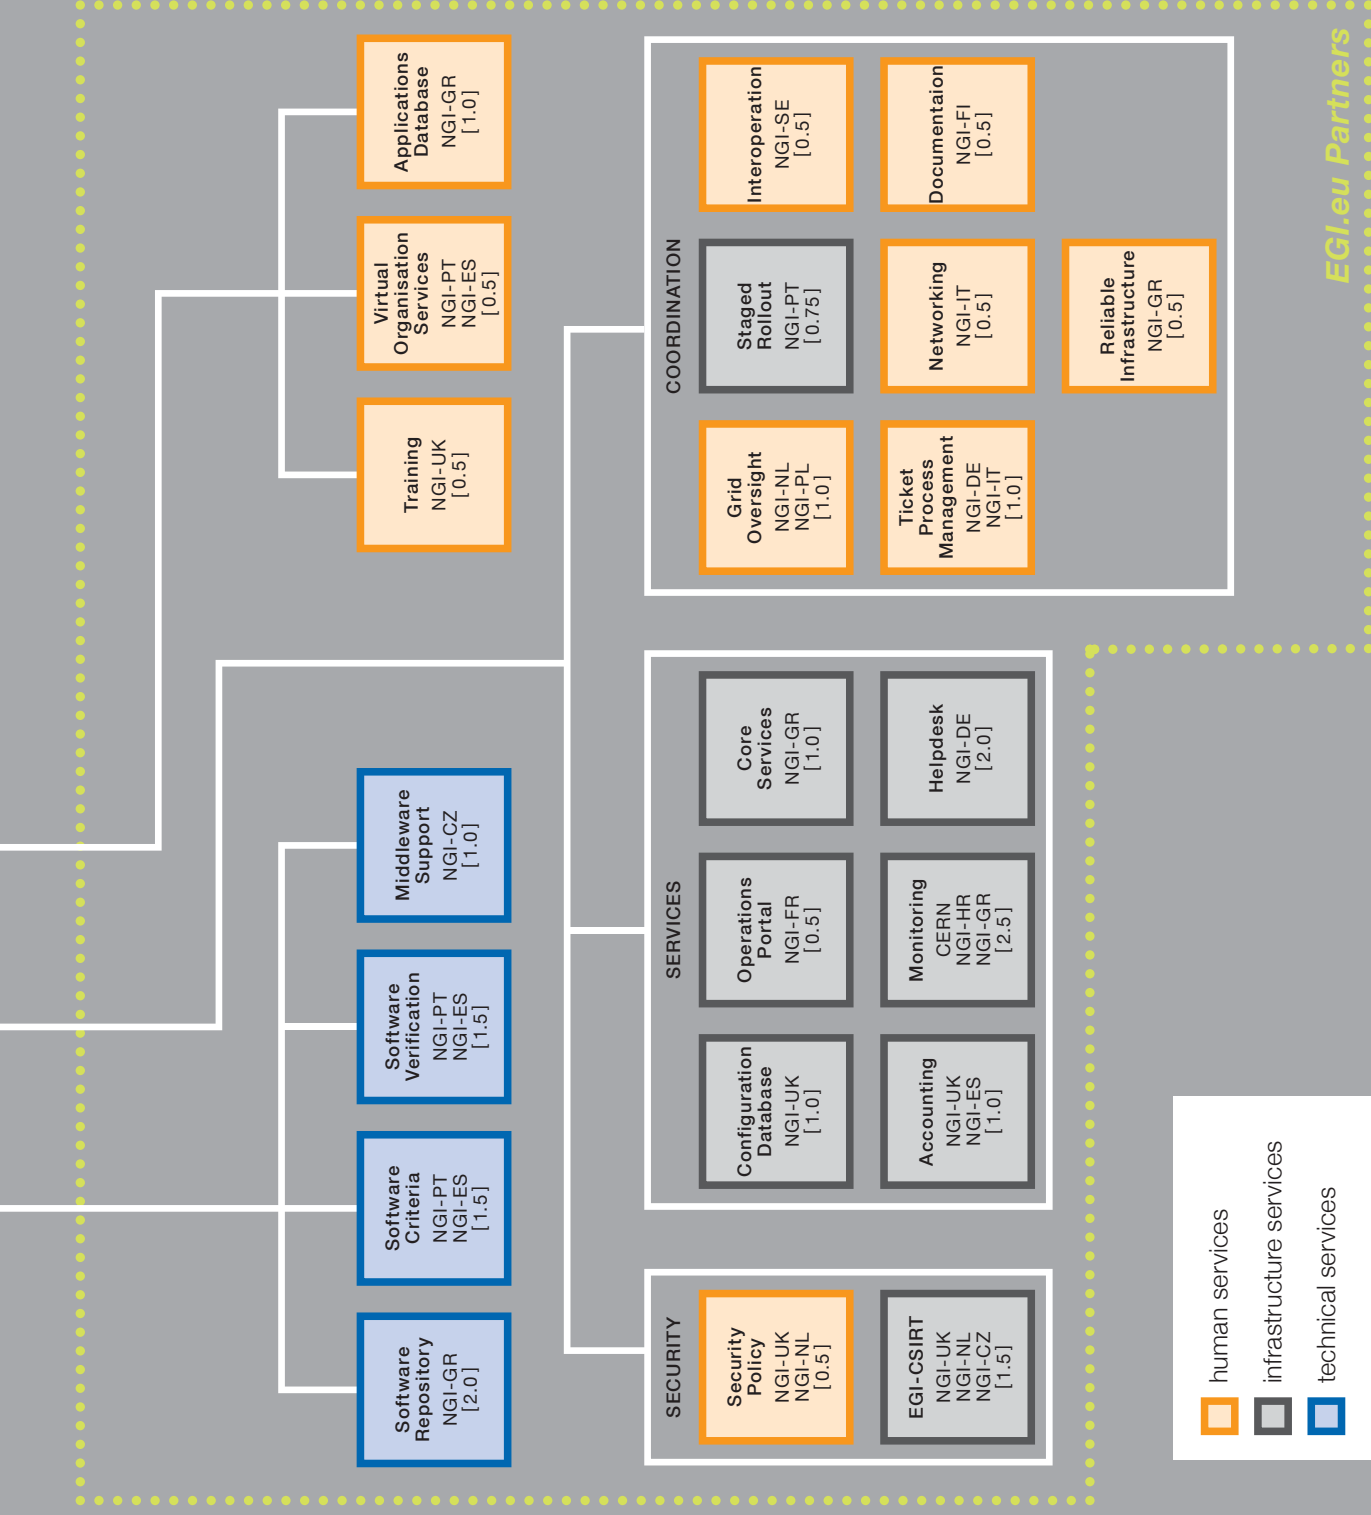
Table showing EGI.eu participation fees for 2011:

PARTICIPANTS	INCOME	PARTICIPANTS	INCOME
NGI-DE	79,880	NGI-RO	19,970
NGI-UK	79,880	NGI-HU	19,970
NGI-FR	79,880	NGI-SK	9,985
NGI-IT	79,880	NGI-HR	9,985
NGI-ES	59,910	NGI-SI	9,985
NGI-NL	39,940	CERN	4,993
NGI-TR	39,940	NGI-CY	4,993
NGI-CH	39,940	NGI-EE	4,993
NGI-BE	39,940	NGI-RU	4,681
NGI-SE	39,940	NGI-BA	3,994
NGI-PL	39,940	NGI-BG	3,994
NGI-NO	39,940	NGI-RS	3,994
NGI-GR	29,955	NGI-LT	3,994
NGI-DK	29,955	NGI-LV	3,994
NGI-FI	29,955	NGI-MK	1,997
NGI-IE	29,955	NGI-ME	1,997
NGI-PT	29,955	NGI-MD	1,405
NGI-IL	19,970	NGI-LU	1,248 [‡]
NGI-CZ	19,970	NGI-AM	281 [‡]
TOTAL		965,128	

Country-based participants pay a participation fee and receive votes based upon their national Gross Domestic Product. Associated participants select their own fee and voting level, subject to the approval of the EGI Council.

EGI.EU MANAGEMENT STRUCTURE





[numbers in square brackets indicate the amount of work effort, shown as full-time equivalent]

CSIRT = computer security incident response team

Glossary

Definitions of key terms are given on this page, together with a list of acronyms and abbreviations used in this document on the opposite page.

NGI codes used in this document:

each National Grid Infrastructure/Initiative is represented by 'NGI-XX', where XX is the two-letter top-level internet domain code for a country.

European Grid Infrastructure (EGI) is a federation of resource providers set up to deliver sustainable, integrated and secure computing services to European researchers and their international partners.

EGI.eu is an organisation based in Amsterdam established to coordinate and manage the infrastructure (EGI) on behalf of its participants: National Grid Initiatives (NGIs) and European Intergovernmental Research Organisations (EIROs).

EGI-InSPIRE (EGI-Integrated Sustainable Pan-European Infrastructure for Research in Europe) is a four-year project, co-funded by the European Commission's 7th Framework Programme (contract number: RI-261323), helping to establish a sustainable, reliable e-infrastructure that can support researchers' needs for large-scale data analysis.

National Grid Initiatives or Infrastructures (NGIs) are organisations set up by individual countries to manage the computing resources they provide to the European e-infrastructure (EGI). NGIs are EGI's main stakeholders, together with two European Intergovernmental Research Organisations (EIROs): CERN and EMBL.

Virtual Organisations (VOs) are groups of people (e.g. scientists, researchers) with common interests and requirements, who need to work collaboratively and/or share resources (e.g. data, software, expertise, CPU, storage space) regardless of geographical location. They join a VO in order to access resources to meet these needs, after agreeing to a set of rules and policies that govern their access and security rights to users, resources and data.

Virtual Research Communities (VRCs) are groups of large-scale research collaborations, or a number of separate VOs grouped according to research domain or computational technique. The group shares information and experience in achieving their goals through the usage of an e-Infrastructure (e.g., best practices, applications, training material)..

Unified Middleware Distribution (UMD) is the integrated set of software components contributed by technology providers and packaged for deployment as production-quality services in EGI.

Photo credits: to be completed after photos are finalised

ATLAS	A Toroidal LHC ApparatuS	HPC	High-Performance Computing
CERN	Européen pour la Recherche Nucléaire (European Organization for Nuclear Research)	HTC	High-Throughput Computing
CLARIN	Common Language Resources and Technology Infrastructure	ICT	Information and Communication Technologies
CPU	Central Processing Unit	IGE	Initiative for Globus in Europe
CSIRT	Computer Security and Incident Response Team	iSGTW	International Science Grid This Week
DARIAH	Digital Research Infrastructure for the Arts and Humanities	IT	Information Technology
EC	European Commission	LHC	Large Hadron Collider
EDG	European DataGrid	MoU	Memorandum of Understanding
EGEE	Enabling Grids for E-science	NGI	National Grid Initiatives or Infrastructures
EGI	European Grid Infrastructure	NIKHEF	National Institute for Subatomic Physics [The Netherlands]
EGI-InSPIRE	EGI-Integrated Sustainable Pan-European Infrastructure for Research in Europe	NIL	NGI International Liaison
EIRO	European Intergovernmental Research Organisation	OGF	Open Grid Forum
EMBL	European Molecular Biology Laboratory	OMB	Operations Management Board
EMI	European Middleware Initiative	PRACE	Partnership for Advanced Computing in Europe
ENVRI	Common Operations of Environmental Research Infrastructures	SIENA	Standards and Interoperability for e-Infrastructure
ERA	European Research Area	SPT	Strategy and Policy Team
ESFRI	European Strategy Forum on Research Infrastructures	TCB	Technology Coordination Board
EU	European Union	UCB	User Community Board
GÉANT	pan-European data network dedicated to the research and education community	UCST	User Community Support Team
GOCDDB	Grid Configuration Database	VAT	Value Added Tax
HEP-SPEC	a benchmark for measuring computer processing power in high energy physics	VO	Virtual Organisation
		VRC	Virtual Research Communities
		VRE	Virtual Research Environment
		VT	Virtual Team



EGl.eu

Science Park 140
1098 XG Amsterdam
The Netherlands

www.egi.eu
contact@egi.eu