

European Grid Infrastructure

2010–11



Contents

Introduction	1
What does EGI.eu do?	3
Role	
Human services	
Infrastructure services	
Technical services	
Impact	10
EGI in the media	
EGI events	
Community	14
Resource infrastructure providers	
Virtual research communities	
Technology providers	
European projects	18
EGI-InSPIRE	
EGI-InSPIRE project results	
e-ScienceTalk	
Future.	24
EGI.eu Financial Statement 2010	26
Balance Sheet	
Income	
Expenditure	
EGI.eu Management Structure	
Glossary.	32

Why e-infrastructures?

Science today is no longer exclusively produced in single research labs or within national boundaries.

Modern scientific challenges call for integrated solutions, increasingly detailed simulations, complex repeatable experiments, cross-country collaborations and, above all, computing power to analyse the vast amounts of generated data. This trend for collaborative discovery will inevitably strengthen in the future and form the basis of exciting innovations in science, society and business. And so will the need for e-infrastructures – geographically distributed computing and storage resources, linked by high-performance networks.

E-infrastructures 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. E-infrastructures are fundamental building blocks of the vision set out by the Digital Agenda and Innovation Europe flagship initiatives, delivered sustainably and as an integrated service across national borders.

Introduction



The ideas behind the European Grid Infrastructure emerged from within the community in late 2006. These ideas were developed through the European Grid Initiative Design Study project and culminated on 8 February 2010 with the coordinating organisation EGI.eu established as a Dutch foundation.

As Chair of the EGI Council, that provides strategic direction to the Community and the Executive Board that supervises the work of EGI.eu, it gives me great pleasure to have seen the establishment of the European Grid Infrastructure and EGI.eu over the last year. The rapid build-up of the organisation and transition of the infrastructure to a federation of National Grid Initiatives is a great achievement by all involved. The structure, activities and plans are described in this the first annual report, and I hope it provides an informative read for the whole community. Although there are many great accomplishments during the first year there are also many challenges ahead. For example, addressing specific needs of different pan-European research communities, integration with e-infrastructures for data and supercomputing, and advancing the service capabilities of the infrastructure.

The success of the first year can be attributed to the creativity of the community, the well-organised NGIs and the support by governments and the European Commission. And especially, on behalf of the EGI Council I would like to thank all the hard working EGI.eu staff and experts in the NGIs.

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



Welcome to the first annual report of the European Grid Infrastructure and its coordinating body EGI.eu.

EGI.eu was established to support the collaborative use of distributed computing resources by European end-users and their international collaborators. In practice, EGI.eu coordinates the European Grid Infrastructure (EGI) on behalf of its participants, the National Grid Initiatives (NGIs) and European Intergovernmental Research Organisations (EIROs). This coordination work started on 1 May 2010.

The first year of EGI.eu, as an organisation based in Amsterdam, was focused on recruiting a team of 20 people to coordinate activities across the operations, user community and technology areas. The Amsterdam team works in close collaboration with partners within the wider EGI community to deliver the technical services and coordination required to run EGI effectively.

From a community perspective the year has seen the production infrastructure complete its transition from a structure built around regional operations centres (ROCs) to a country-based structure. EGI.eu has also established many of the human, technical and infrastructure services and procedures needed for EGI to operate securely and effectively across Europe.

My thanks to all of those involved in and around EGI.eu for your support during this first year.

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**

- | | | |
|--------------------------------|-------------------------------|----------------------------------|
| Belgium: NGI-BE, BELNET | Greece: NGI-GR, GRNET | Norway: NGI-NO, UNINETT Sigma AS |
| Bulgaria: NGI-BG, IPP-BAS | Hungary: NGI-HU, NIIF | Poland: NGI-PL, CYFRONET AGH |
| CERN | Ireland: NGI-IE, Grid-Ireland | Portugal: NGI-PT, UMIC |
| Croatia: NGI-HR, SRCE | Israel: NGI-IL, IUCC | Romania: NGI-RO, ICI |
| Cyprus: NGI-CY, CyGrid | Italy: NGI-IT, INFN | Serbia: NGI-RS, IPB |
| Czech Republic: NGI-CZ, CESNET | Latvia: NGI-LV, LUMII | Slovakia: NGI-SK, SlovakGrid |
| Denmark: NGI-DK, DCSC | Lithuania: NGI-LT, VU | Slovenia: NGI-SI, ARNES |
| Estonia: NGI-EE, EENet | Luxembourg: NGI-LU, RESTENA | Spain: NGI-ES, CSIC |
| Finland: NGI-FI, CSC | Macedonia: NGI-MK, MARGI | Sweden: NGI-SE, SNIC |
| France: NGI-FR, CNRS | Montenegro: NGI-ME, UoM | Switzerland: NGI-CH, SWITCH |
| Germany: NGI-DE, Gauß-Allianz | The Netherlands: NGI-NL, NCF | Turkey: NGI-TR, ULAKBIM |
| | | United Kingdom: NGI-UK, JISC |

**in alphabetical order by [country, where applicable]: [abbreviation used in this document], [short name of the lead organisation]

What does EGI.eu do?

ROLE

EGI.eu's role is to coordinate the resources distributed across Europe to deliver an integrated and secure e-infrastructure for the benefit of its multi-disciplinary users. EGI.eu works with user communities to build a 'virtuous cycle' of development, gathering user requirements and collaborating with technology providers to feed them into the production infrastructure.

EGI.eu coordinates human, infrastructure and technical services within the community for the use of the whole community. Human services refer to the coordination and support work needed at a European level to integrate the different aspects of the community. Infrastructure services provide integration to the production infrastructure, which represents the resources offered to our user community to support their research activities. Finally, the technical services are crucial to help the distributed community to work more effectively together.

These services are currently supported by the participants in EGI.eu through the fees they pay, direct support from the partners who host these services, and the European Commission through the EGI-InSPIRE project. Details on these services are provided in the following pages.

Governance

The EGI Council is responsible for defining the strategic direction of EGI and the EGI collaboration. The Council acts as the senior decision-making and supervisory authority of EGI.eu. The Council participants are the National Grid Initiatives (NGIs) that represent each country and European Intergovernmental Research Organisations (EIROs). The Council delegates oversight of the day-to-day running of EGI.eu to the seven members of the Executive Board. In turn, the EGI.eu Executive Board devolves financial and organisational responsibility to the Director.

NGIs and EIROS

The European NGIs participate in the Council and also in the EGI-InSPIRE project, most but not all in both. NGIs and EIROs, such as CERN, provide a wide range of technical services to support the operation of a national e-infrastructure that is integrated with the European Grid Infrastructure. In each country, the NGI is the single point of contact to all the institutions and research communities related to a national grid infrastructure.

The NGI in each member state needs to provide these services so that it can interact with EGI and support its user communities. The activities of the NGIs are not limited to the tasks that each NGI performs at a national level to maintain its own infrastructure, but also include global and international tasks that allow IT resources to be shared in a robust and transparent way across the whole community.

HUMAN SERVICES

Human services, and the coordination they provide, are a vital component of the complex EGI ecosystem. The managerial, technical, operational, policy and dissemination activities provided by EGI.eu in Amsterdam are a resource-efficient integrated way to deliver key services to the whole EGI community.

Management

The EGI.eu secretariat provides support to the Council, Executive Board and the EGI collaboration as a whole. The Council delegates the technical governance of EGI to a number of boards, which are also maintained and

supported by EGI.eu staff. The Operations Management Board works to ensure that the production infrastructure operates effectively and securely. The User Community Board provides a forum to discuss and prioritise requirements as they emerge from the end-users. These requirements are fed into the Technical Coordination Board, which balances the needs of end-users, operations and the capabilities offered by the external Technology Providers.

Support the user community

The EGI.eu User Community Support Team (UCST) coordinates NGI support activities around Europe to continuously meet the evolving needs of the distributed

user communities. The UCST helps new and established user communities to make the most out of the resources provided by EGI and the supporting technical services.

Coordinated operations

EGI.eu coordinates and supervises the operational activities of the production infrastructure with the support of the community. Assignment and escalation of support tickets are handled by the ticket process management teams (at NGI-DE and NGI-IT). General oversight of the infrastructure (handled by NGI-PL and NGI-NL) is supported by networking expertise (by NGI-IT), the validation and distribution of monthly availability and reliability statistics (by NGI-GR), interoperation between the different middlewares within and external to NGI (by NGI-SE) and operational documentation (by NGI-FI). The EGI Computer Security and Incident Response Team (from NGI-UK and NGI-NL) provides coordination security and training to improve the response to incidents (e.g. security drills) and security vulnerabilities.

External relations

The EGI.eu Policy Development Team (PDT) coordinates policy-development across the various EGI boards and governance bodies. Policies can be 'technical' through open standards, 'operational' to steer operational procedures, or 'strategic' to provide analysis and guidance around key topics of discussion. The PDT has prepared key strategic reports around EGI's sustainability plans, integration of clouds and virtualisation into EGI, EGI's role within the Europe 2020 strategy, the applicability of the ERIC legal framework for EGI.eu and an analysis of the standards relevant to EGI. Specialist security policy expertise (provided by NGI-UK and NGI-NL) promotes

EGI IN NUMBERS (1 MAY 2011)

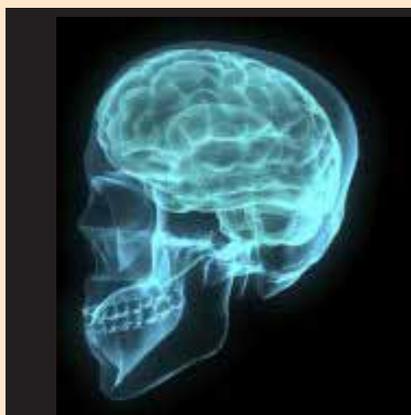
- 33 EGI Council participants
- 50 EGI-InSPIRE project partners
- 41 resource providers
- 11,226 users and 219 Dissemination teams
- CPU cores (EGI-InSPIRE partners): 239,840
- Computing capacity (EGI-InSPIRE partners): 1,894,861 HEP-SPEC 06
- Usage: 26.6 million jobs/month*
- Reliability (as defined below): 92.73%*
- Availability (as defined below): 93.85%*

*monthly average taken from 2010–2011 figures

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.

DIAGNOSING ALZHEIMER'S



The World Health Organisation estimates that about 18 million people worldwide have Alzheimer's disease, a figure projected to double by 2025. Patients with early symptoms have trouble remembering recent events. As the brain degenerates further, they lose mental and physical functions and need 24-hour care. Doctors do not know what causes Alzheimer's.

Key to treating Alzheimer's is an early diagnosis of the condition. Early diagnostic tools are available based on brain scan images, but doctors need to compare patient scans to a large amount of healthy patient data – the more data, the more accurate the diagnosis. This is where the grid can help.

DECIDE (Diagnostic Enhancement of Confidence by an International Distributed Environment) is setting up a dedicated grid infrastructure that brings diagnostic tools to the medical community. Over the grid, doctors can compare patient data with a federated database of images located all over Europe, gathered into a single virtual repository. Diagnostic tools such as SPM (Statistical Parametric Mapping) can then track the progression of the disease over time, from mild cognitive impairment through to the later stages of the disease.

www.eu-decide.eu

integrated identity and trust policies within EGI and partner infrastructures.

The EGI.eu dissemination team coordinates the NGI dissemination teams around Europe to communicate the work of EGI within the community, as well as to decision-makers and the wider audience. Two large conferences each year guarantee the continued building and strengthening of the EGI community.

INFRASTRUCTURE SERVICES

Infrastructure services are the key to keep EGI's production infrastructure secure and running 24/7.

Configuration database

EGI relies on a central configuration database (managed by NGI-UK) to record static information contributed by the resource providers as to the service instances that they are running and the individual contact, role and status information for those responsible for particular services (<http://goc.egi.eu>). This database is a source of information for many other operational tools, automated mailing lists, and other services that enable EGI to operate efficiently at a European scale.

Help is there

The EGI Helpdesk (run by NGI-DE) is at the heart of the EGI's support structure to end-users and operations staff and can be integrated with ticketing systems used within the NGIs. The helpdesk (<http://helpdesk.egi.eu>) allows issues, submitted via a ticketing system, to be allocated to support units in the NGIs or in VRCs. The users responsible for submitting the tickets can also track their progress and resolution.

Monitoring the infrastructure

EGI's Monitoring Infrastructure is a distributed message-based service (maintained by NGI-HR and NGI-GR) where each NGI regularly tests the level of functionality delivered by services in their production infrastructure. The results from the tests are used to generate automatic alarms and

tickets in case of critical failures and to compute monthly availability and reliability statistics. The central monitoring services (maintained by CERN) collect and store the NGI test results in databases that can be accessed through the MyEGI portal (<http://grid-monitoring.cern.ch/myegi>). These services are key to ensure that resources are always up and running reliably when users wish to use them.

Accounting for its usage

The central accounting service in EGI (provided by NGI-UK) includes a repository to store accounting records and a portal (<http://accounting.egi.eu>) to visualise the accounting information (provided by NGI-ES). The central databases are populated through single or aggregated usage records published by the Resource Centres.

Operations portal

EGI has a centralised portal (developed by NGI-FR) to help the operations community to run the infrastructure effectively (<http://operations-portal.egi.eu>). The portal supports grid oversight activity by providing capabilities such as the broadcast tool to announce upcoming infrastructure changes, Dissemination team (VO) information and management facilities to see which VOs are enabled on which sites, or a dashboard for the operations community to display information about failing monitoring probes and related tickets.

Staged rollout

New software updates are deployed into the production infrastructure incrementally through a staged rollout to ensure that they are reliable in actual use. The staged rollout process (coordinated by NGI-PT) tests new software through an early deployment by expert resource

centres, following successful verification of the software component against published criteria.

A secure infrastructure

The EGI Computer Security and Incident Response Team (EGI CSIRT) works to provide a secure infrastructure for EGI's users and communities. Their coordination efforts are supported by tools and services that identify vulnerabilities where they appear on resource centres through the use of out of date software. One of these services is the Pakiti tool (maintained by NGI-CZ), which performs random software configuration checks on individual resource centres.

Core services

EGI (through NGI-GR) also provides core services that are generally under the responsibility of individual VOs, for its own use and for the benefit of smaller VOs that do not yet have the critical mass to develop their own instances. Examples of core services include top-level BDII and Workload Management Services, the catch-all Certificate Authority, central file catalogues or the VOMS service for VO membership management.

BIG SCIENCE NEEDS BIG COMPUTING



It's no mean feat to accelerate particles around a 27km ring close to the speed of light, then colliding them 800 million times a second. Only one machine in the world is capable of this job – the Large Hadron Collider (LHC), a particle accelerator located 100m underneath the French/Swiss border.

The unique nature of the LHC attracts global interest and the 8,000 particle physicists who work on its experiments are from almost every nation on the planet. The sheer size of data generated, 15 petabytes each year, and the geographically distributed nature of the researchers

means that a worldwide LHC computing grid (WLCG, a global collaboration linking grid infrastructures and computer centres worldwide) is the most suitable computing solution. EGI is a major partner in the WLCG, contributing over a billion hours of CPU time to the project in the last 12 months. EGI was working closely with physicists from the very beginning (even before the machine was switched on in 2008) to help integrate grid services with the unique demands of the LHC.

<http://public.web.cern.ch/public/en/LHC/LHC-en.html>

TECHNICAL SERVICES

EGI's technical services aim to improve the usage of the production infrastructure and to support the delivery of the technology that makes up the production infrastructure. Many of the centrally-provided services can be embedded within NGI or VO web portals for the benefit of their specific communities.

Dissemination team services

The VO Services team (NGI-PT and NGI-ES) currently provides a monitoring system that user communities can customise with tests specific to their community. These VO services also include:

1. guidance and support for the start-up, management and operation of VOs (https://wiki.egi.eu/wiki/VO_Services);
2. advice about tools, services, documentation and procedures to optimise resources usage, and
3. VO evaluation and supporting documentation to user communities who wish to operate VO services themselves. The EGI VO Services portfolio will be expanded as the needs of new and established communities grow.

Software services

EGI relies on external technology providers to deliver software to meet the requirements of end-user and operations communities, defined in the Unified Middleware Distribution roadmap. Technology providers contribute their software components to a centralised Software Repository (managed by NGI-GR). Every component uploaded to the repository (<http://repository.egi.eu>) is validated against defined quality criteria (by NGI-ES and

NGI-PT) before being placed into staged rollout. This activity is supported through a number of customised tools (provided by NGI-CZ). The acceptance criteria (developed by NGI-ES and NGI-PT) are based on the prioritised requirements defined by the Technical Coordination Board. The criteria are public and capture the key functional and non-functional features expected from the delivered technologies (<http://go.egi.eu/QualityCriteria-1>).

Applications Database – AppDB

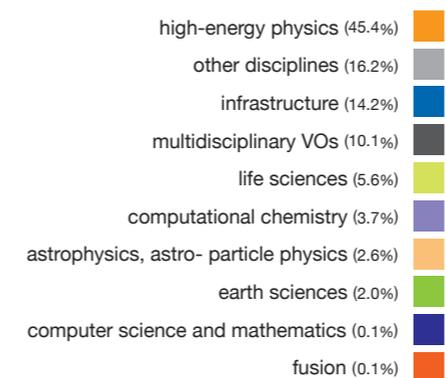
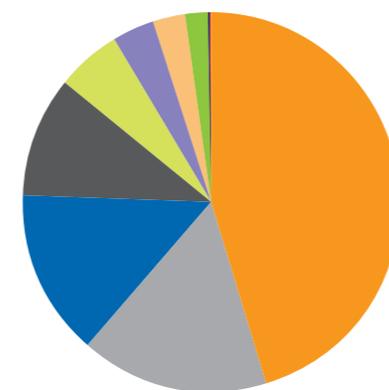
The EGI Applications Database (developed by NGI-GR) stores tailor-made computing tools for scientists to use (<http://appdb.egi.eu>). It embraces all scientific fields, from resources that simulate exotic excitation modes in physics, to applications for complex protein sequences analysis. Storing pre-made applications and tools means that scientists do not have to spend research time developing their own software. The database's goals are:

1. to inspire scientists less familiar with programming to use EGI and its resources due to the immediate availability of the software that they need to use; and
2. to avoid duplication of effort across the user community.

Training marketplace

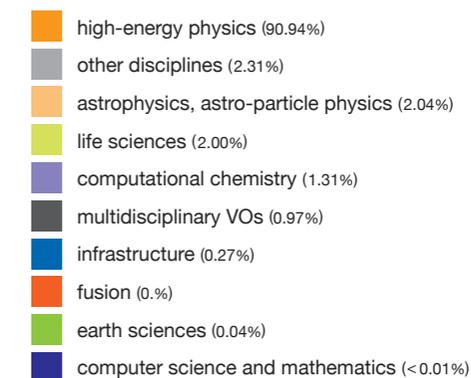
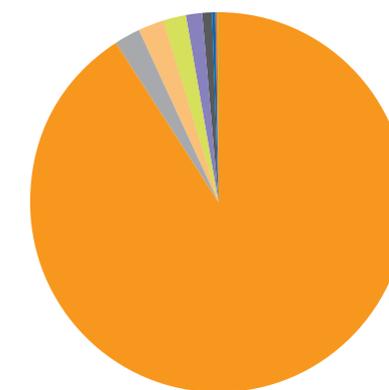
The training marketplace (managed by NGI-UK) connects trainers and users in the NGIs and VRCs by enabling them to exchange information relating to training materials and training events between different communities in different localities within EGI (<http://training.egi.eu>).

Registered grid users by research disciplines:



Total number of users = 13,319 (March 2011)

Grid computing demand by research disciplines:



Total = 6,633,607,020 HEP-SPEC 06 (monthly accumulative total, from May 2010–April 2011)

RESOURCE ALLOCATION

While EGI is a federation of resources coordinated by EGI.eu, there is no central allocation or control of who uses these resources. Resource providers choose which users they support, and the users determine which dissemination teams they are a part of. If the

user's research collaborations could be strengthened by sharing their resources with international collaborators within their dissemination team, then this can be enabled by the resource provider. Through its services, EGI therefore enables a collaborative network for computing and data resources that can underpin researcher communities' social networks.

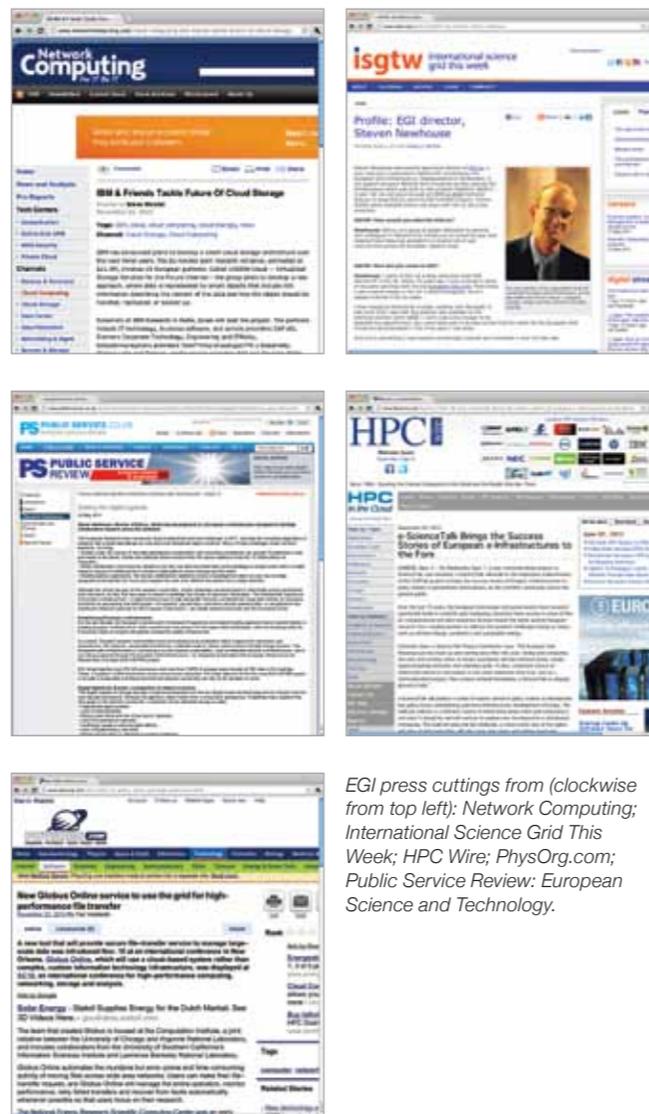
Impact

EGI IN THE MEDIA

The EGI.eu dissemination team is on hand to answer any queries from journalists about the organisation, its mission and goals. The team puts journalists in contact with scientists working on research projects using the grid infrastructure, arranges interviews with EGI managers, Executive Board and Council members, and advises freelance journalists with ideas for stories.

As well as dealing with news as it happens, there is a concentrated press effort during EGI.eu's two annual flagship events – the User Forum and the Technical Forum. The team identifies stories derived from the forums that would be of interest to general or specialised press, from innovative developments of the grid to interesting scientific results obtained using the grid.

EGI.eu issued two press releases for the Technical Forum in September 2010: a media invitation to a press conference and an announcement of the start of the project. "EGI-InSPIRE brings together European e-infrastructure community" was issued to 3,870 journalists through the AlphaGalileo press service and was also published on the Cordis news wire, the EGI website and sent to the media contacts list and the dissemination mailing list. A press release announcing the funding for EGI-InSPIRE was also issued by the EC Press Office. The initiative led to 27 press cuttings, including articles in HPCwire, iSGTW, ZDNet, Yahoo News, ITnews in Australia, Science Business and Environment & Energy Management. The local Dutch media were targeted in partnership with NGI-NL, and press interviews were held with Web Wereld and Tweakers.net.



EGI press cuttings from (clockwise from top left): Network Computing; International Science Grid This Week; HPC Wire; PhysOrg.com; Public Service Review; European Science and Technology.

EGI.eu dissemination team also assists partners and collaborators with coordinated press activities, such as issuing joint press releases to enhance media interest. For example, as a result of the EGI User Forum in April 2011, a press release was written in collaboration with the Lithuanian technology provider BAIP, highlighting the innovation grid technologies brought to health care and government services in Lithuania.

EGI.eu is engaged in a range of social media channels, with associated RSS news feeds, to complement traditional press outreach. These media connect the organisation directly with interested and targeted audiences:

- EGI Blog (www.egi.eu/blog)
- Twitter (http://twitter.com/egi_inspire)
- You Tube (www.youtube.com/europeangrid)

HEARING THE ANCIENT WORLD

The ancient world still 'talks' to us through the artefacts and stories left behind; but one thing we rarely truly experience is ancient sound. Recreating an instrument that has not been used for over a thousand years is not a simple task. Not only have the skills been lost to the mists of time, a single instrument is also difficult to share with the entire archaeological community.

Now the ASTRA (Ancient instruments Sound/Timbre Reconstruction Application) project is changing that. Using various sources, such as artefact images and descriptions in texts, the Parma-based team has recreated the music of the ancient world. Knowing the shape and the materials used in the original instrument, they can simulate its characteristics and the tones it would create.

With a high-powered PC taking 4 hours to recreate just 30 seconds of sound, they realised they needed more resources. Using the European grid they have successfully recreated four different ancient instruments from Europe and South America. The result is a set of files that sound like the original instrument and can be played on electronic equipment by anyone. They have even been used in live performances alongside modern instruments.

www.astraproject.org

EGI EVENTS

EGI Technical Forum 2010

The EGI Technical Forum 2010 was held in Amsterdam at the Beurs van Berlage from 14 to 17 September 2010 in partnership with the BiG Grid project, the Dutch NGI. This was the first major event within the EGI community and brought together European distributed computing projects and their collaborators in academia and business, from around Europe and around the world. The major theme of the meeting, achieved through technical sessions, a demonstration and exhibition area, networking space and events, was to establish collaborations between the new and the current European Distributed Computing Infrastructure projects to meet the needs and requirements of the research community.

More than 570 delegates registered for the EGI Technical Forum, of which 17% were female. The event included 290 contributions in the form of presentations, demos, posters and workshops.

EGI User Forum 2011

The EGI User Forum 2011 was organised by EGI.eu, Vilnius University and LITNET in Vilnius, Lithuania, 11-14 April 2011, with the support of the EGI-InSPIRE and European Middleware Initiative (EMI) projects and local secretariat BAIP. The conference took place at the Radisson Blu Lietuva and showcased the diversity of the user community within EGI through plenaries, oral presentations, poster sessions and co-located workshops. The programme also included numerous networking and opportunities to 'meet the experts'.

In total, 427 delegates registered for the event, of which 18% were female. The event featured 196 contributions, 173 speakers and 34 session conveners. Over 250 images were uploaded to the Flickr photostream and more than 30 posts were added to the GridCast blog about the event from 8 bloggers, including 14 videos and slide shows. Over 2,600 unique visitors visited the main event website, representing 20,000 page views.



EGI User Forum 2011 – Kostas Glinos, Head of the 'GÉANT & e-Infrastructure' unit at the European Commission, talking to Steven Newhouse and Alberto Di Meglio from EMI (European Middleware Initiative) (above left); preparation for a presentation on the DECIDE project, held at the EGI exhibition booth (below left); User Forum poster (above right).



EGI Technical Forum 2010 – a presentation at the forum in Amsterdam (below right); postcard to promote the forum ahead of the event (centre).



Community

EGL.eu has been established by the EGL community, is governed by representatives of the EGL community and serves the needs of the EGL community. As part of this role, EGL.eu is helping to structure the community by formalising the relationship it has with the different elements within it, including the e-Infrastructure Reflection Group (e-IRG), the Open Grid Forum (OGF) and the SIENA initiative.

Other working relationships established within EGL's first year, internal and external to Europe, are detailed in the following sections.

RESOURCE INFRASTRUCTURE PROVIDERS

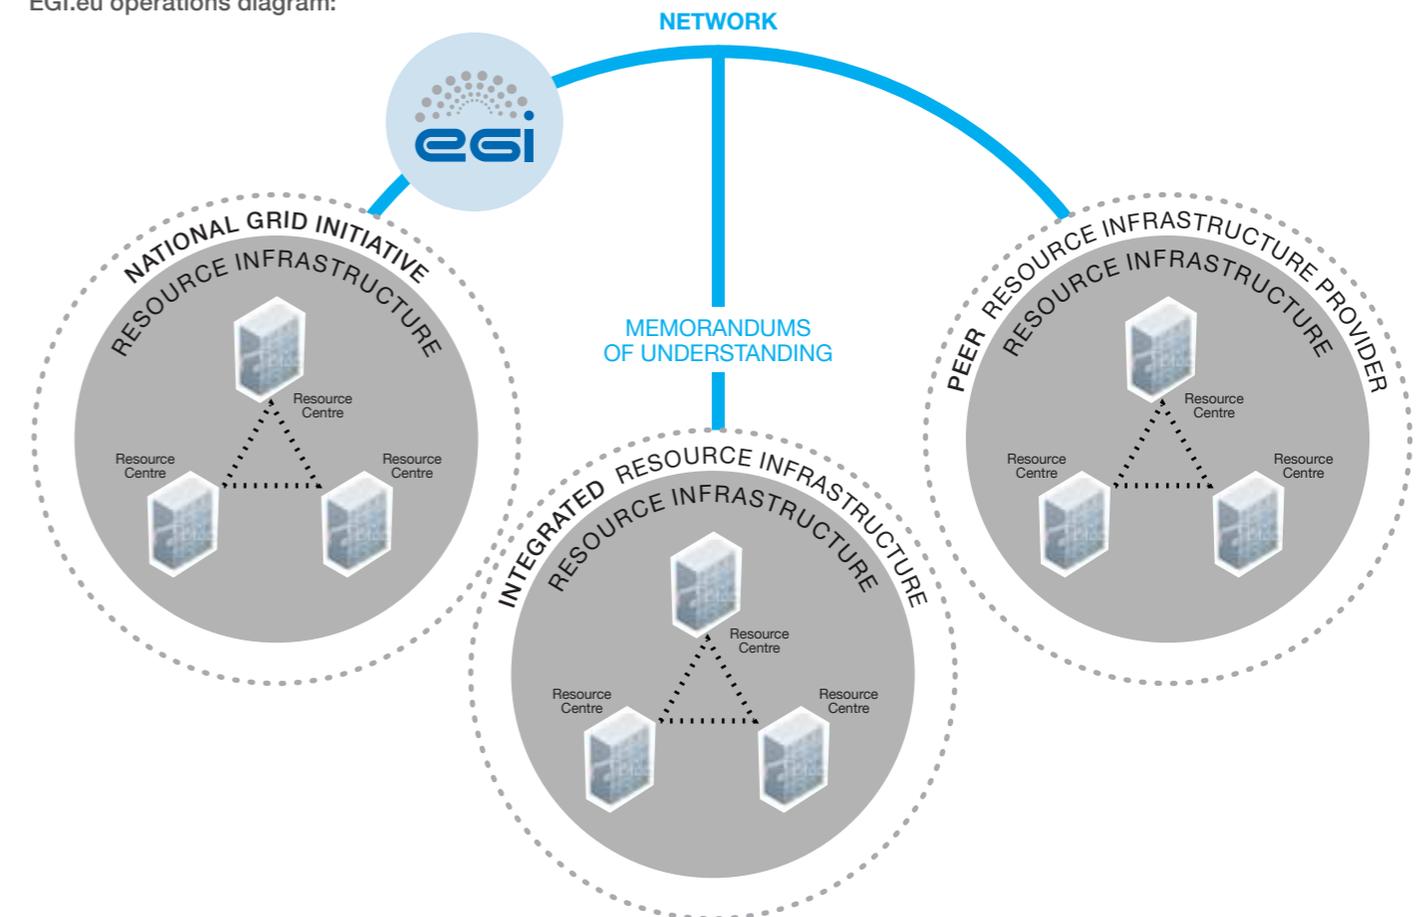
Resource Infrastructures are federations of Resource Centres created to deliver integrated services to their stakeholders and managed by Resource Infrastructure Providers, such as the NGIs and EIROs that participate in EGL.eu. Each of these organisations provides resources and services, to their geographical (e.g. national or regional) or domain-specific (e.g. High Energy Physics, Life Sciences) user communities. EGL.eu's Resource Infrastructure Providers are represented in the Operations Management Board (OMB) that develops and defines the technical procedures required to federate the different Resource Infrastructures reliably.

All Resource Infrastructures must be operated in a consistent similar way to ensure that users can access resources irrespective of geographical location or Resource Infrastructure Provider. EGL.eu provides the infrastructure and human coordination services that ensure seamless access to a site's functional services throughout

Europe. EGL.eu is able to deliver an infrastructure that can support a variety of diverse user communities by ensuring that the functional services are working correctly at all of the sites, by offering a helpdesk system for tracking when they are not, and by accounting the resource usage by users or communities. The expected availability and reliability of the resources are defined by Operational Level Agreements between EGL.eu and the Resource Infrastructure Providers.

EGL.eu is also working to establish relationships with e-infrastructures based throughout the world to enable European researchers and scientists to collaborate and share resources with colleagues outside Europe. Strong relationships already exist with the GISELA project in Latin America, and the Open Science Grid in North America to support the High Energy Physics community. Eight countries in the Asia Pacific area are partners within the EGL-InSPIRE project.

EGL.eu operations diagram:



GISELA

GISELA brings together 19 beneficiaries (and 12 other parties) from 15 countries in Europe and Latin America for a two-year project partly funded by the European Commission. It started in September 2010 with two main goals:

- To implement a Latin American Grid Initiative (LGI) sustainability model rooted on National Grid Initiatives (NGI) or Equivalent Domestic Grid Structures (EDGS), in association with CLARA, Latin American NRENs and EGL.

- To provide Virtual Research Communities with the e-infrastructure and application-related services required to improve the effectiveness of their research.

The long-term idea is to establish a powerful e-infrastructure facility built on the legacy of the 'e-science grid facility for Europe and Latin America' (EELA) project series, as well as to develop and implement a sustainability model for the e-infrastructure.

VIRTUAL RESEARCH COMMUNITIES

EGL's user community has two important structures: Dissemination teams (VOs) and Virtual Research Communities (VRCs). A VO is a group 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 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). A VRC can be a group of large-scale research collaboration, or a number of separate VOs grouped according to research domain or computational technique. The VRC model allows communities to interact with EGL (through defined points of contact) across wider areas. VRCs are supported by EGL through technical services, such as the applications database, the training marketplace and VO Services.

EGL.eu works with VRCs representative of multiple scientific fields to guarantee that the infrastructure evolves in a direction that is truly user-driven. VRCs are represented within EGL at the User Community Board (UCB). The UCB provides a route for communication between user communities and infrastructure providers on topics such as requirements, dissemination, events or plans for the future.

EGL also provides the mechanisms needed to link resources contributed by individual user communities in a secure, integrated infrastructure for access by the relevant dissemination teams.

TECHNOLOGY PROVIDERS

Requirements gathered within the operations and Virtual Research Communities are analysed and prioritised through the Technology Coordination Board (TCB). The TCB works with external technology providers—EMI, IGE, SAGA and StratusLab (see page 17)—to define requirements and acceptance criteria (in terms of reliability, scalability and functionality) for the software that they deliver to EGL. Building on its needs, EGL produces the EGL Technology Roadmap describing how it expects the production infrastructure to evolve over time. The Unified Middleware Distribution (UMD) Roadmap, a subset of the EGL Technology Roadmap, describes the components being source from within the EGL community. To improve EGL's sustainability the number of technology components coming from just within the EGL community needs to reduce over time.

WeNMR

WeNMR is a VRC that brings together research teams in structural biology into a worldwide virtual research community, focusing on biomolecular Nuclear Magnetic Resonance (NMR) and Small Angle X-ray Scattering (SAXS). WeNMR provides these communities with virtual platforms and user-friendly computation tools, supported by an underlying high performance e-infrastructure.

EUROPEAN MIDDLEWARE INITIATIVE

The European Middleware Initiative (EMI) is a collaboration of four major European middleware providers: ARC, dCache, gLite and UNICORE. The initiative aims to deliver a consolidated set of middleware components for deployment in EGL, as part of the Unified Middleware Distribution. EMI also provides middleware to PRACE and other distributed computing infrastructures.

www.eu-emi.eu



SAGA PROJECT

The SAGA project (Simple API for Grid Applications) provides a programming abstraction that encapsulates and hides from the application developer the differences between middleware systems. The ultimate aim of SAGA is to make distributed computing simple, easy, reliable and robust. SAGA software is developed through an open source model and has recently become an OGF (Open Grid Forum)-approved standard.

<http://saga.cct.lsu.edu>



INITIATIVE FOR GLOBUS IN EUROPE

The Initiative for Globus in Europe (IGE) aims to be a comprehensive service provider for the European e-infrastructures regarding the development, customisation, provisioning, support (including training), and maintenance collaboration with EGL, other distributed computing infrastructure projects, and standard development organisations.

www.ige-project.eu



STRATUSLAB

StratusLab is a project set up to develop a complete, open-source cloud distribution that allows grid and non-grid resource centres to offer and to exploit an Infrastructure as a Service cloud. The project focuses on enhancing distributed computing infrastructures such as EGL and uses agile methodologies to rapidly evolve the distribution to meet cloud user and administrator demands. StratusLab operates a reference infrastructure to get feedback from users and to validate the distribution.

<http://stratuslab.eu>



European projects

EGI-INSPIRE

The goal of the EGI-InSPIRE project is to enable the transition from a project-based system to a sustainable pan-European e-infrastructure. This is done by supporting 'grids' of high-performance computing (HPC) and high-throughput computing (HTC) resources to evolve into federations of national and domain-specific Resource Infrastructure Providers.

EGI-InSPIRE also supports the integration of Distributed Computing Infrastructures (DCIs) such as clouds, supercomputing networks and desktop grids, to benefit user communities within the European research area. EGI.eu, the coordinator of the EGI-InSPIRE project, contributes to this effort with the support provided by its policy, dissemination, operational, user community and technical teams. Additional work within the project oversees the development of the operational tools and the current heavy

users of the infrastructure, such as high energy physics, computational chemistry, materials science, earth science, fusion, astronomy & astrophysics and life sciences, as they move their critical services and tools from a centralised support model to one driven by their own individual communities.

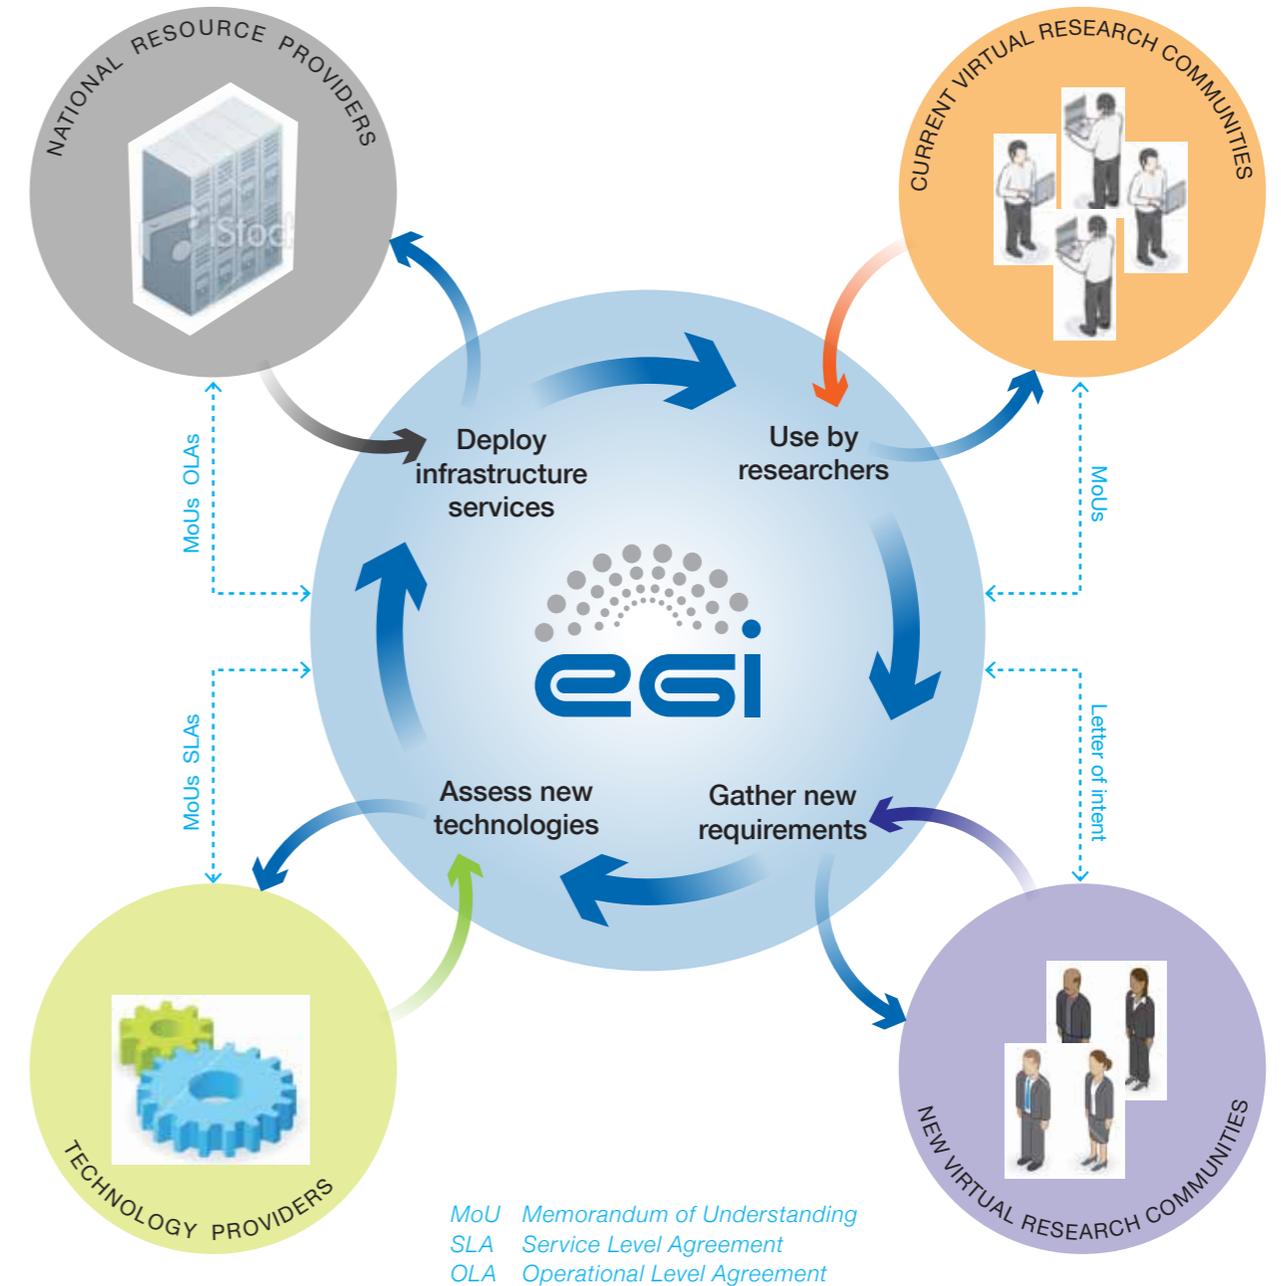
The project focuses on these activities through a series of work packages spread across networking, service and joint research activities.

The EGI-InSPIRE supported services drive a virtuous feedback cycle – starting with a set of integrated services on the production infrastructure that meets user needs, working with external technology providers to define new or improved services based on these existing services, assessing the quality of the new delivered services, followed by their deployment in to the production infrastructure.

EGI-INSPIRE IN NUMBERS

- 4-year project, May 2010 to April 2014
- €25 million contribution from the EC
- Project cost €72 million
- Total financial contribution around €330 million
- Total effort contributed 9,261 person months
- Effort equivalent to 193 people a year
- 50 project partners
- 38 National Grid Initiatives
- 2 European Intergovernmental Research Organisations
- 9 unfunded Asia Pacific partners

The EGI-InSPIRE virtuous feedback cycle:



EGI-INSPIRE PROJECT RESULTS

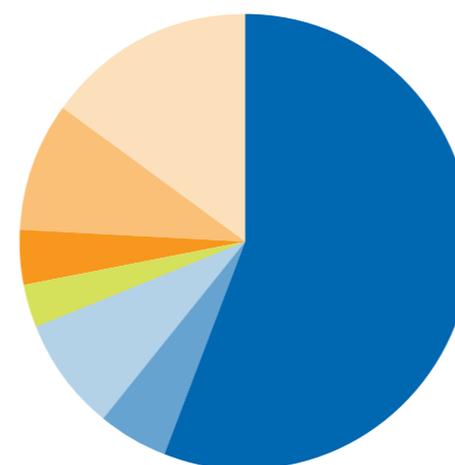
EGI-InSPIRE has had a productive first of its four-year project in establishing a sustainable secure production infrastructure for multi-disciplinary use for the European Research Area.

At the centre of the sustainable EGI model is EGI.eu, now fully staffed with 21 employees in Amsterdam, which has established the governance and technical structures for an open European wide e-infrastructure. Many of the regional operations centres established within EGEE have now transitioned to national operations centres supporting 38 NGIs and their user communities. The structuring provided by virtual research communities is allowing EGI to evolve the centrally supported services it provides and promote software reuse between communities.

The user and operations communities have published more than 70 papers and conference proceedings. These have included papers in journals such as Journal of Physical Chemistry, Physics in Medicine and Biology, Journal of Grid Computing, Earth Science Informatics, Grid and Cloud Database Management and the Journal of Physics: Conference Series. In addition, the Book of Abstracts from the EGI User Forum in Vilnius in April 2011 included over 150 abstracts from the oral presentations.

EGI's activities are promoted its own events (such as the User and Technical Forums) and throughout the community by events such as Ibergrid 2010 in Braga Portugal, the GridKa school in Karlsruhe Germany, a grid day at the University of Helwan in Egypt, the Cracow Grid Workshop, the NGS Innovation Forum in Oxfordshire UK, Scientific Gateways and Visualisation at the European Geosciences Union 2010 and the 8th e-infrastructure Concertation Meeting at CERN in Geneva.

EGI-InSPIRE – distribution of project effort between activity (total 9,192 person months):



- SA1 Operations (56%)
- SA2 Provisioning the software infrastructure (5%)
- SA3 Services for heavy user communities (8%)
- JRA1 Operational tools (3%)
- NA1 Project management and administration (4%)
- NA2 External relations (9%)
- NA3 User community coordination (15%)

KEY RESULTS

Project Overview

- EGI-InSPIRE Presentation <http://go.egi.eu/43>
- EGI-InSPIRE Paper <http://go.egi.eu/201>

Community

- EGI Technical Forum, Amsterdam, Netherlands www.egi.eu/EGITF2010
- EGI User Forum, Vilnius, Lithuania <http://uf2011.egi.eu>
- EGI User Forum 2011 Book of Abstracts <http://go.egi.eu/EGIUF2011BoA> (PDF file 2.6MB)
- Work of the Asia Pacific Region <http://go.egi.eu/316>

Technical Architecture

- Roadmap for Interactions with Other DCI Projects <http://go.egi.eu/207>
- EGI Operations Architecture <http://go.egi.eu/218>
- Integration of Clouds and Virtualisation into the European production infrastructure <http://go.egi.eu/258>
- User Community Support Process <http://go.egi.eu/61>
- Alignment of EGI.eu with the ERIC organisational model <http://go.egi.eu/244>
- Capabilities Offered by the Heavy User Communities to Other Communities <http://go.egi.eu/154>

EGI Annual Activity Reports

- External Relations <http://go.egi.eu/377>
- User Community <http://go.egi.eu/386>
- Production Infrastructure <http://go.egi.eu/413>
- Software Provisioning activity and the work of DMSU <http://go.egi.eu/375>
- Heavy User Communities Tools and Services <http://go.egi.eu/312>
- Operational Tool Maintenance and Development <http://go.egi.eu/372>



Above left: the GridCafé website.
Below left: a copy of the e-ScienceBriefings, an issue on supercomputing.



Above right: the e-ScienceTalk website (www.e-sciencetalk.org). Centre: video interview with Rimantas Žylius, the Minister of Economy for Lithuania, at the EGI User Forum for the GridCast blog. Below right: the Physics Island at the e-ScienceCity



E-SCIENCETALK

'e-ScienceTalk: Talking about e-science'

The e-ScienceTalk project started on 1 September 2010 and brings the success stories of Europe's e-infrastructure to policy makers in government and business, to the scientific community and to the general public. e-ScienceTalk is coordinated by EGI.eu and works with the European Grid Infrastructure and collaborating projects to report on the interactions of grid computing with e-infrastructures, such as cloud computing and supercomputing.

The project identified a strong need for reporting targeted at influential policy makers in government, science and business. The e-ScienceBriefings interpret EC policy documents and reports in an accessible and attractive format, illustrating scientific results and impacts. This year, e-ScienceBriefings have been published giving an overview of supercomputing, cloud computing and the e-infrastructure landscape. e-ScienceTalk also hosted the 150-strong 8th e-Infrastructure Concertation meeting at CERN in November 2010.

The GridCafé website was launched by CERN in 2003 to explain to the public "what grid computing is and what it could soon be." e-ScienceTalk aims to keep the GridCafé website at the cutting edge of e-Science dissemination and is adding new areas to create an online e-ScienceCity and a 3D virtual GridCafé. Team members and scientists blog live from events around the world on our GridCast site, which includes videocasts and slide shows. The GridGuide website gives a human face to

e-Science, allowing users to listen to podcasts from grid sites worldwide and read interviews with researchers. The GridGuide is closely linked to the Real Time Monitor, a 3D visualisation that shows data traffic moving around the grid in real time.

International Science Grid This Week (iSGTW) is a free weekly online newsletter that promotes grid computing and e-Science around the world by sharing stories of science and scientific discoveries. Now reaching over 8,000 subscribers in more than 195 countries and territories, it is an international publication with a European editor based at CERN, a US editor at Fermilab, as well as several contributing writers from around the world including Asia, Latin America and Africa. Recently relaunched with a new interactive website, iSGTW now offers readers the opportunity to comment and rate articles, take part in polls, publish blogs, share stories on social media sites and announce their events.

e-ScienceTalk is working with other projects such as GÉANT, DANTE and DEISA/PRACE to disseminate the interdependencies of Europe's e-infrastructures through the e-ScienceBriefings, articles in iSGTW and by expanding the content of the GridCafé. e-ScienceTalk is also working closely with e-IRG and the European Strategy Forum on Research Infrastructures (ESFRI) projects who are currently involved in building a united roadmap for the development of e-infrastructures in Europe. e-ScienceTalk has recently signed MoUs with e-IRGSP2/3, EUIndiaGrid2, LinkSCEEM-2, WeNMR, CHAIN, DEGISCO and EMI and is looking forward to working closely with these projects in the future.

Future

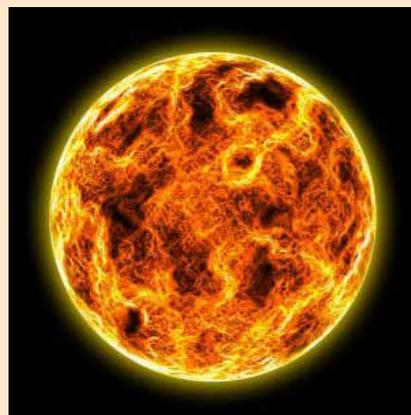
New technology will continue to drive the evolution of the services that EGI offers to its user communities. The next year will see the full integration of Globus and UNICORE into the production infrastructure to meet the needs of user communities using HPC resources. Alongside the existing resources already offered by EGI, such as computing and storage, virtualised resources will start to become available. Virtualised resources should allow end-user communities to have greater control over the software environments that are deployed on their behalf across the infrastructure.

The outreach to new user communities and the support given to existing communities will continue. The evolving service infrastructure will attract user communities that have not been able to benefit from proceeding projects service offerings. Customisable web gadgets for all of EGI's key services will be made available to our resource providers and virtual research communities for integration into their own support portals. Expanding the range of services – deployed both centrally and locally – that can benefit user communities will be achieved by promoting the work taking place in individual NGIs to the whole community.

While continuing to deliver and evolve the production infrastructure and the services offered to user communities around it, EGI.eu will continue to host discussions on sustainability and business models. EGI.eu, either directly in Amsterdam or through partners in the community, delivers services to NGIs, EIROs and virtual research communities (as described earlier). These services are provided for EGI's federation of resource providers and virtual research communities (a B2B relationship) to support site administrators or researchers (a B2C relationship).

European e-infrastructures have been recognised over the last decade as providing a foundation for innovation across the European Research Area. They are also recognised as a key component in Europe's strategy for 2020 through the Innovation Union and the Digital Agenda for Europe. Sustainability, stability and integration of these e-infrastructures are therefore critical. Structures such as the European Research Infrastructure Consortium (ERIC) offer one model of an organisation that federates national activities for the governance, funding and seamless service delivery across Europe. The migration of EGI to such a body will be explored in the years to come.

BUILDING A SUN ON EARTH



One of the biggest problems confronting the planet is energy. Between dwindling fossil fuel reserves and a desire to move to clean sources of power, a solution is needed soon. One possible alternative technology is the process that powers our sun – fusion.

The International Thermonuclear Experimental Reactor (ITER) aims to demonstrate the feasibility of fusion power. But this requires computing power that pushes the limits of state of the art resources. To meet such requirements, the EUFORIA (EU Fusion FOR Iter Applications) project has been

setup to provide a comprehensive framework and infrastructure for modelling and simulating the reactor. The EGI is one of the main partners, contributing over a million hours of CPU time in the last year

www.euforia-project.eu/EUFORIA/

MAPPING A COUNTRY'S GENOME



Since the publication of the complete human genome in 2003, there has been huge interest in the benefits genomics can bring to individuals. National governments are interested because studying the genetic variation in their populations might help develop new medical treatments and diagnostic techniques.

One such project is the Genome of the Netherlands (GoNL). The researchers want to sequence the genomes of 750 individuals to create a snapshot of the Dutch genetic profile. The input data alone is 30 terabytes, with the output expected to be 20 times larger.

The analysis is very computationally intensive – a standard desktop PC would require a run time of about 13 years.

In 2010, the GoNL project approached EGI's Dutch partner BiGGrid for help. Since then, work has moved quickly with initial trials showing analysis times being cut by 80%. The next step is to this scale up, which works for a subset of the data, into a production service that can process the 30 terabytes of data as and when the researchers need it.

www.bbmri.nl/en-gb/activities/projects/131-genome-of-the-netherlands

EGL.eu Financial Statement 2010

The following accounts relate from the establishment of EGL.eu on 8 February 2010 to 31 December 2010. The complete financial accounts and accompanying statement are available online (<http://go.egi.eu/568>). The accounts were adopted by the EGL Council on 14 April 2011.

BALANCE SHEET

ASSETS		€
Trade and other receivables ¹	16,016	
Cash and cash equivalents ²	1,457,386	
	1,473,402	
LIABILITIES		€
Reserve acceptable costs ³	561,967	
Current liabilities ⁴	911,435	
	1,473,402	

Figures as of 31 December 2010

¹ Debtors (€1,940), interest (€7,579), receivables from employees (€6,497)

² EGL.eu current account (€850,516), EGL.eu Savings Account (€606,296), Other (€574)

³ Reserve for E-tasks that have not yet been undertaken (EGL.eu: €103,470 and partners €41,930) and for EGL.eu (€416,567)

⁴ Taxes and social securities (€59,924), project pre-financing (€471,792), EGL-InSPIRE money to paid to partners (€231), accounts payable (€44,992), E-Task to partners (€275,851), accountant (€8,330), administration (€2,185), other amounts payable (€378), and accrued staff pay (€47,756).

INCOME

INCOME 2010		€
Start of year balance	468,150	
INCOME		
EGL InSPIRE ⁵	354,060	
eScienceTalk	30,964	
EGL.eu participants	1,000,000	
Interest	7,603	
Total income	1,392,626	
TOTAL	1,860,777	

⁵ The income from the EC for the EGL-InSPIRE project is based upon the provisional cost statement (i.e. effort and travel) and may change due to audit and project amendments.

EXPENDITURE

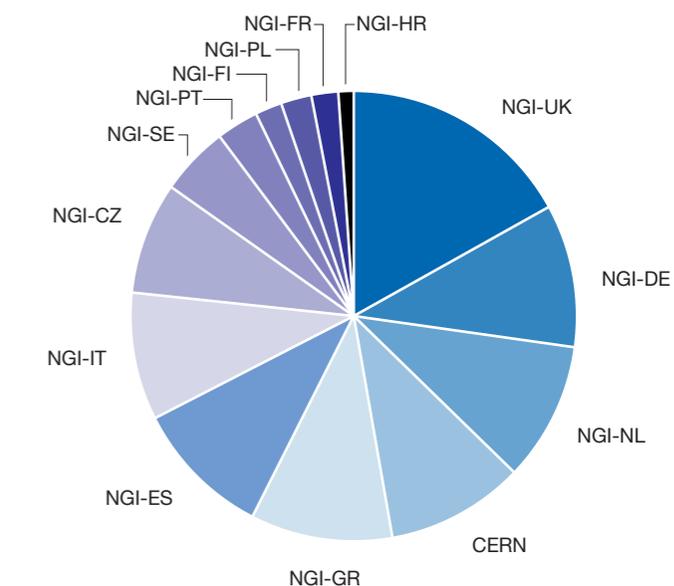
EXPENDITURE 2010		€
Staff salaries ⁶	690,087	
Subsidy for EGL Global Tasks	398,436	
Office costs ⁷	150,620	
Direct project costs ⁸	59,668	
TOTAL	1,298,810	

⁶ EGL.eu was recruiting staff during 2010 and only delivered 104PM over 8 months of operation, instead of the 150PM expected over the same period. Around 10% of the cost relates to non-recurring staff recruitment or relocation costs. The staff salaries are determined through the EGL.eu Terms of Employment (<https://documents.egi.eu/document/269>), local tax and labour laws, and the contracts offered to individual members of staff.

⁷ Includes travel undertaken by staff. EGL.eu benefited from support by the local host (NGI-NL and the BiGGrid project) during the start-up phase by paying for the office space, some of the non-recurring professional fees and other startup costs.

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

Subsidy paid by EGL.eu for the delivery of EGL Global Task:



NGI-UK	€67,571	NGI-ES	€38,501	NGI-FI	€9,328
NGI-DE	€41,661	NGI-IT	€37,416	NGI-PL	€8,362
NGI-NL	€39,606	NGI-CZ	€30,839	NGI-FR	€7,864
CERN	€39,247	NGI-SE	€19,907	NGI-HR	€4,618
NGI-GR	€39,062	NGI-PT	€13,260		

Table showing EGI.eu participation fees for 2010:

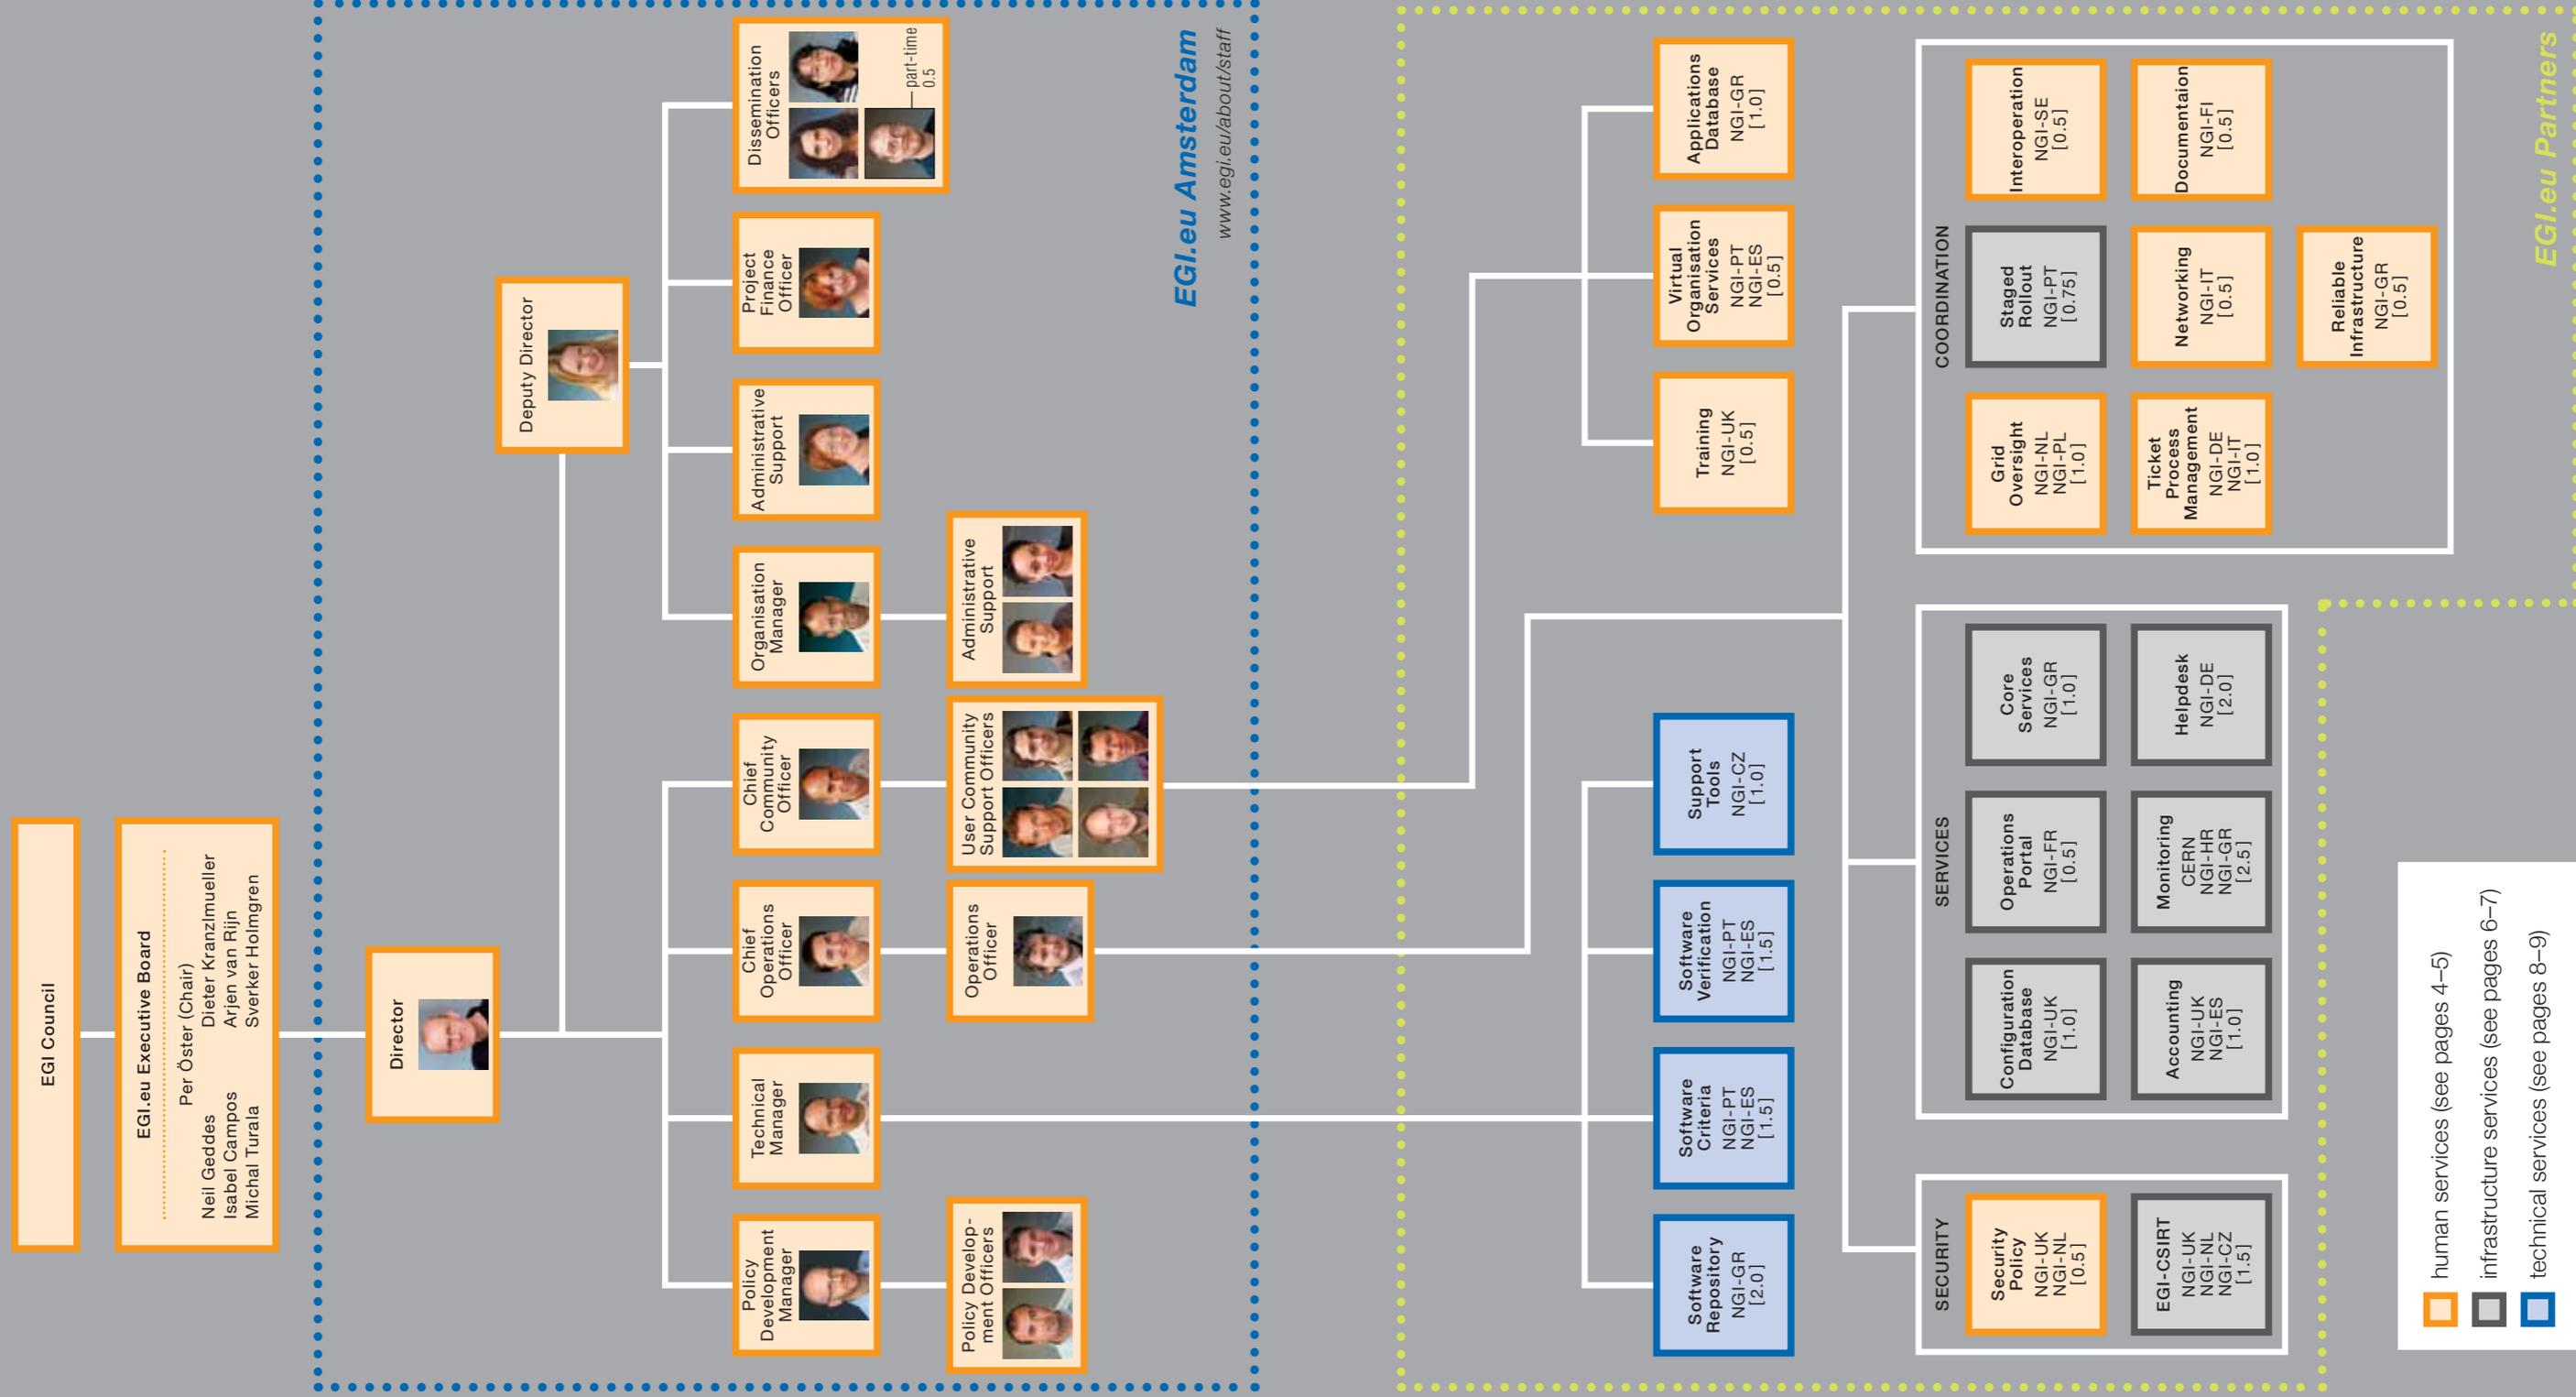
PARTICIPANTS	INCOME €	PARTICIPANTS	INCOME €	PARTICIPANTS	INCOME €
NGI-DE	79,880	NGI-NO	39,940	NGI-SI	9,985
NGI-UK	79,880	NGI-GR	29,955	CERN	4,993
NGI-FR	79,880	NGI-DK	29,955	EMBL	4,993
NGI-IT	79,880	NGI-FI	29,955	NGI-CY	4,993
NGI-ES	59,910	NGI-IE	29,955	NGI-EE	4,993
NGI-AT	39,940	NGI-PT	29,955	NGI-BG	3,994
NGI-NL	39,940	NGI-IL	19,970	NGI-RS	3,994
NGI-TR	39,940	NGI-CZ	19,970	NGI-LT	3,994
NGI-CH	39,940	NGI-RO	19,970	NGI-LV	3,994
NGI-BE	39,940	NGI-HU	19,970	NGI-MK	1,997
NGI-SE	39,940	NGI-SK	9,985	NGI-ME	1,997
NGI-PL	39,940	NGI-HR	9,985	NGI-AL	1,498
TOTAL					1,000,000

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.

EMBL = European Molecular Biology Laboratory

EGI.EU MANAGEMENT STRUCTURE



EGI.eu Amsterdam
www.egi.eu/about/staff

- human services (see pages 4–5)
- infrastructure services (see pages 6–7)
- technical services (see pages 8–9)

[numbers in square brackets indicate the amount of work effort, shown as full-time equivalent]
CSIRT = computer security incident response team

Glossary

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).

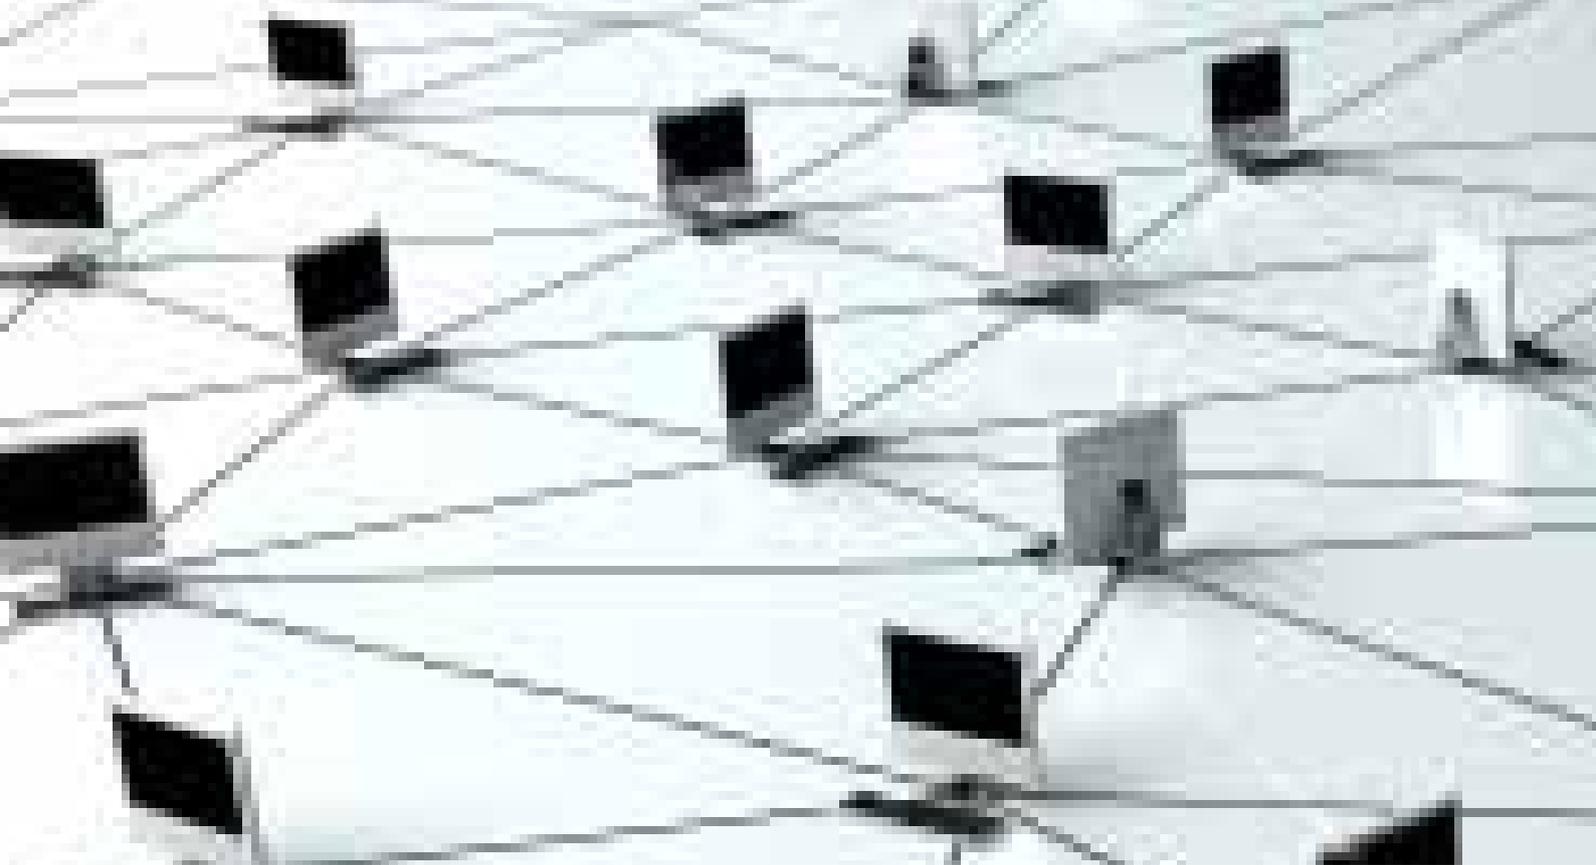
EGI-InSPIRE (EGI-Integrated Sustainable Pan-European Infrastructure for Research in Europe) is a four-year project, co-funded by the European Commission's 7th Framework Programme (contract number: RI-261323), helping to establish a sustainable, reliable e-infrastructure that can support researchers'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

Dissemination teams (VO) represent the individuals with on-line access to resources that span multiple administrative and geographical boundaries.

Virtual Research Communities (VRC) are collaborations of 'like-minded' individuals grouped together by discipline or computational model. VRCs can span multiple VOs.

Unified Middleware Distribution (UMD) is a roadmap that provides EGI's vision of the long-term development of the software deployed on the infrastructure. The roadmap describes the features that are necessary to operate a distributed and federated computing infrastructure.



EGl.eu

Science Park 140
1098 XG Amsterdam
The Netherlands

www.egi.eu
contact@egi.eu

PRINTER TO INSERT
FSC LOGO HERE

EG11005