Inspired

Spring 2011

European Grid Infrastructure

News from the EGI community

This Issue

Spring at *Inspired* is all about the EGI User Forum in Vilnius held between 11-14 April. Viviane Li closes the curtain on the forum and Sy Holsinger tells us how the policy workshop went. Neasan O'Neill reviews the release of EMI 1 and I bring you a summary of the MoUs and SLAs signed at the User Forum. But that's not all...

We also have a review of the EGI operations architecture and latest infrastructure figures as well as a profile of the GISELA project set up to develop a sustainable e-Infrastructure in Latin America.

Simon Hettrick, from the Software Sustainability Institute, shares his ideas on how to foster research collaborations and Dimitra Kotsokali, from SEERA-EI, highlights the conclusions of a report on e-Infrastructures.

If you want to contribute with ideas, suggestions or stories to the newsletter don't forget to let me know!

Sara Coelho
sara.coelho@egi.eu

The First Swallows (pirmosios kregzdės) herald springtime in Vilnius, Lithuania (photo: Neasan O'Neill)

EMI: Scaling new heights

Neasan O'Neill on the launch of EMI1

One of the main problems facing the distributed computing infrastructure community is interoperability. The various disciplines, requirements and resources mean that the solutions that have been created may not always work together.

The European Middleware Initiative (EMI) is starting to solve this problem and last week saw their first full release, EMI-1.

Funded for three years by the European Commission, EMI is bringing together middleware experts from around Europe, and further afield. However, they know they have a difficult task and in honour of this have decided to name their releases after some of the biggest mountains in Europe. EMI 1 is codenamed Kebnekaise, the double peaked mountain that is the highest in Sweden.

While there is a mountain to climb EMI are not doing it alone. They are working with the existing solutions, ARC, gLite and UNICORE, including contributions from these three, as well as other technology providers, like dCache. The collaboration hopes to create a product to help improve access to all distributed computing resources.

Currently the future looks bright for EMI. Having signed a Memorandum of Understanding and a Service Level Agreement with EGI earlier this year, starting official discussions with PRACE and now their first full release the project hopes to move onwards and upwards.

Florida Estrella, the project's Deputy Director, says: “The team has worked extremely hard over the last year, we are sure this is the first step to creating a unified middleware for distributed computing infrastructures. We hope that future releases see us really benefit the users and providers throughout Europe and beyond.” •

More Information

http://www.eu-emi.eu/emi-1 (including release notes for the software)

User Forum 2011 wrap up

Viviane Li closes the curtain

Sun sets over Vilnius as the User Forum drew to a close on 14 April (photo: Corentin Chevalier)

In the crisp cool air of a late-sprung Baltic springtime, more than 400 participants gathered in Vilnius, Lithuania over a four-day period to attend the EGI User Forum 2011. At the Radisson Blu on the edge of the Old Town, six conference tracks hosted a diverse programme of presentations, demonstrations, exhibitions, posters and workshops.

Kostas Glinos, Head of the 'GÉANT & e-Infrastructure' Unit at the European Commission, commenced proceedings to an overfilled auditorium at the opening plenary and outlined the European vision for e-Infrastructure. Steven Newhouse, director of EGI.eu, followed with a summary of the last year's achievements by the EGI-InSPIRE project and Alberto Di Meglio, director of the European Middleware Initiative (EMI), described the latest initiatives to deliver sustainable open source middleware stacks for European researchers.

This year's User Forum, held in conjunction with the EMI Technical Conference, attracted a mixture of grid users and software developers. Six separate tracks reflected this diversity, from 'User Environment' to 'Technologies for Distributed Computing'.

"The 'User Support Services' track was really successful," enthused Gergely Sipos, a member of EGI.eu's User Community Support Team. "We presented our services, discussed tools and the available training. The more users are aware of what we can offer, the better we can serve them. Also, all the services we have now, including those that previously existed in EGEE [Enabling Grid for E-sciencE], have been restructured and are better integrated with the EGI website."

Another important outcome of the Forum for Sipos was the connection with new Virtual Research Communities (VRCs). "We made contact with VRCs that we weren't familiar with before. We can now follow up with these contacts and try to understand their communities' requirements better, so they can be met by the technology providers."

End-users and software developers attended a lively exchange at the 'Meet the Experts' session hosted by EMI. The session aimed to connect users and developers with technology providers to create better middleware.

From a local perspective, the Forum helped to reinforce the current development of computing infrastructure in Lithuania and to empower local scientists. Algimantas Juozapavičius from Vilnius University, project leader of the Lithuanian National Grid Infrastructure, is encouraged by the productive exchanges between computing specialists, the EGI policy team and local governmental officials. "This [distributed computing] infrastructure gives more resources in a more convenient way. Also, it stimulates researchers to expand their research and encourage additional contact with Lithuanian researchers."

As the event drew to a close, and after the hefty hearty zeppelins (Lithuanian potato dumplings) which may have landed in your stomachs have been digested, EGI.eu is proud to have hosted its first User Forum to great success.

"The breadth and diversity of the EGI Community continues to amaze me," said Newhouse. "The efforts spent over the last year defining what we can offer to user communities and how we can interact with them have been clearly demonstrated during the User Forum.

Highlights

> Rimantas Žylius, Minister of Economy for Lithuania, welcomed the EGI community to Vilnius
> EMI launched EMI-1
> EDGI launched MetaJob
> 10 MoUs and SLAs signed

Winners:

> Best demonstration: 'xGUS – a helpdesk template', Sabine Reisser (with Torsten Antoni)
> Best exhibition: StratusLab
> Best poster: 'Environmental Protection and Modeling Applications in Bulgaria – Deployment and Results', Todor Gurov et al
> Domenico Vicinanza won the Forum's Treasure Hunt event

For your diary:

> User Virtualisation Workshop: 12–13 May 2011, Amsterdam, the Netherlands
> EGI Technical Forum 2011: 19–23 September, Lyon, France

MoUs and SLAs: collaborations from theory to practice

Sara Coelho was at the EGI stand to report on the agreements

EGI.eu aims to provide European researchers with an integrated, reliable and sustainable e-infrastructure to support data-intensive research. This ambitious goal depends on a solid relationship with partners outside the EGI community, but equally committed to the development of e-Infrastructures. Fostering and developing strong working relations with such partners was one of the main objectives of EGI.eu's first year of operations. Many months of discussions followed, identifying the common goals and mutual interests of EGI.eu and its ever-growing number of external partners.

The User Forum in Vilnius saw the culmination of this work with a flurry of signings of Memoranda of Understanding (MoUs) and Service Level Agreements (SLAs) with technical and infrastructure providers, as well as virtual research communities. The venue was the EGI stand in the exhibition area and the host was Steven Newhouse, director of EGI.eu.

"One of the biggest achievements of EGI's first year of operation has been the collaborations established with software providers via the Technical Coordination Board and the structuring of virtual research communities being represented within the User Community Board," says Newhouse. "We now have MoUs signed with a variety of middleware providers and user communities. For the most critical middleware components we have SLAs in place that will provide a predictable response to any support requests coming from the EGI community."

So far, EGI has established MoUs and SLAs with the European Middleware Initiative (EMI) and the Initiative for Globus in Europe (IGE). The SAGA project (Simple API for Grid Applications), which provides a programming abstraction that encapsulates the differences between middleware systems, became the third technical provider to enter a collaboration with EGI.

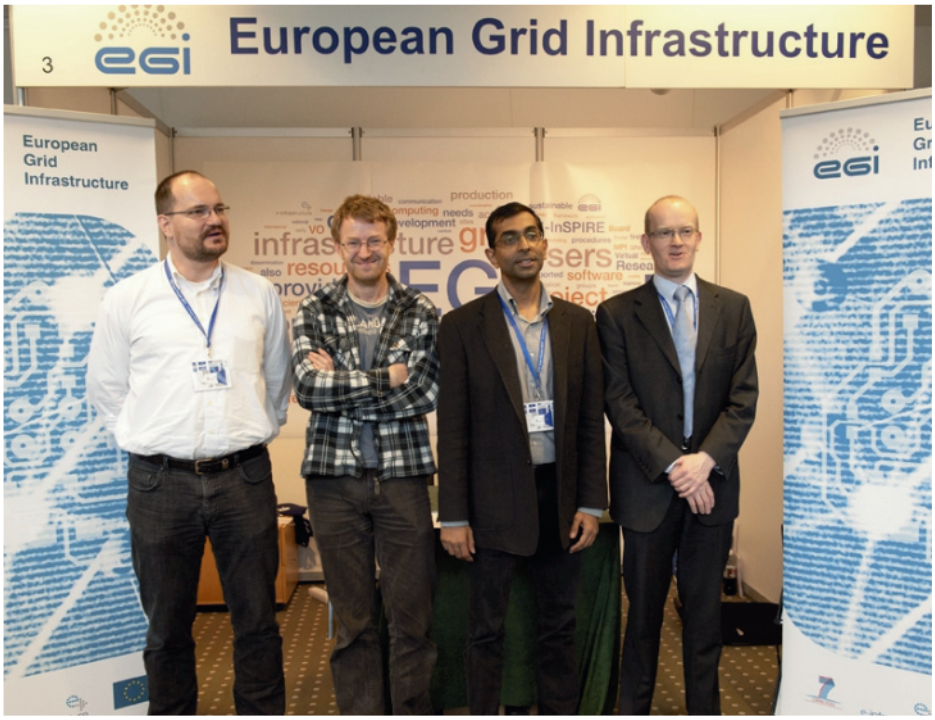

From left: Michel Drescher, SAGA's André Merzky & Shantenu Jha and Steven Newhouse at the EGI booth in the Vilnius User Forum.

"The ultimate aim of the [SAGA] project is to make distributed computing simple, easy, reliable and robust," said Shantenu Jha, SAGA project leader, in an interview to the GridCast blog. The SAGA software has recently become an OGF-approved standard and has started to be adopted by several e-infrastructure projects from all over the world.

"It's a well-known fact that EGI is the largest and probably the most important open source grid infrastructure in the world," said Jha in Vilnius. "Being able to work with EGI is not only important for us in terms of advancing impactful, cutting-edge science, but is also a great way to stay involved with the grid community."

While MoUs are useful to define a common ground and a general course of action, SLAs have an added value – they clarify how the business relationship is going to work.

(cont)

EGI.eu has signed SLAs with:
> European Middleware Initiative www.eu-emi.eu
> Initiative for Globus in Europe www.ige-project.eu
> SAGA project http://saga.cct.lsu.edu

... and MoUs with:
> GISELA consortium www.gisela-grid.eu
> WeNMR www.wenmr.eu
> StratusLab http://stratuslab.org
> SAGA

e-ScienceTalk signed MoUs with:
> GISELA consortium
> WeNMR
> CHAIN project www.chain-project.eu

"The SLAs with SAGA and IGE allow us to support new grid usage models and to broaden the range of middleware we can provide to our user communities," says Michel Drescher, EGI.eu's Technical Manager. "And the SLA signed with EMI is an important milestone securing the current day-to-day operational business."

The quest to provide a wide range of technical solutions to its user community has set EGI on a common route with StratusLab, a two-year FP7 project. StratusLab was set up to develop open source software that allows users to run customised virtual environments inside a cloud created and managed by resource centres.

"We also want to show that this can be done with grid services," says Charles Loomis, StratusLab's project director. Loomis ran a successful workshop in the EGI User Forum which saw more than 20 users setting up virtual machines in a cloud operated with StratusLab software.

Left: Steven Newhouse shakes hands with Charles Loomis, from StratusLab; Catherine Gater and Alexandre Bonvin sign an MoU between the e-ScienceTalk project and the We-NMR infrastructure.

In the long term, the StratusLab aims to provide stable and robust cloud software for grid services. "EGI sites are one of our target communities and key to the sustainability of the StratusLab distribution," says Loomis. GRNET already has two EGI-certified sites operated as StratusLab clouds.

"The experimental work that is being undertaken within the StratusLab project provides an important data point on how EGI as a whole may evolve," says Newhouse. "This is an area we will explore at the EGI User Virtualisation workshop on the 12-13 May in Amsterdam". •

EGI User Virtualisation Workshop

Sy Holsinger unveils the plan

Virtualisation and cloud computing have demonstrated how new technologies can enable dynamic execution environments or on-demand elastic service deployment with new, clear cost measurements and business models. Today's ICT policies and services that are tailored to the established e-Infrastructure user communities do not always meet the needs of newcomers. Now is the time for EGI to provide a more flexible, efficient e-Infrastructure to attract new users on a larger scale.

EGI is organising a dedicated workshop on User Virtualisation (12-13 May 2011, Amsterdam) to bring together three critical groups within the European production infrastructure – resource providers, end-users and technology providers. The workshop will address specific questions through a series of user specific presentations and topical breakouts and serves as a follow-up to a recently published report on the integration of clouds and virtualisation into EGI (http://go.egi.eu/258). A technical roadmap defining the EGI cloud profile is also currently underway.

The workshop will deliberately not spend time on technology related presentations or provide an introduction to cloud computing, as there have been plenty of recent activities in the area (e.g. DCI collaborative and SIENA roadmaps) as well as recently held events (e.g. Cloudscape-III).

Discussions will be focused on understanding if and how EGI should move towards providing an Infrastructure as a Service (IaaS) model to support data intensive research communities and if this service should be provided by federations of resource providers from the research community. Other key topics to be covered include what sort of use (consuming communities and use cases) would the research community make of such a capability and what are the major issues (e.g. technical, policy, governance, etc.) that need to be resolved for this use to happen.

The outcome of the meeting will be a set of critical observations that can be developed into a roadmap that will meet the growing need for virtualised resources from the European research communities. •

For more information or to register, as places are limited, visit: http://go.egi.eu/uvw1

Project profile: GISELA

Plotting a sustainable e-Infrastructure in Latin America

The GISELA consortium – short for Grid Initiatives for e-Science virtual communities in Europe and Latin America – has many things in common with EGI, including a mission to provide sustainable e-Infrastructure resources to its community of scientists and researchers.

It's no surprise then that project coordinator Bernard Marechal joined Steven Newhouse in the User Forum in Vilnius to sign two Memoranda of Understanding (MoUs): one project MoU between GISELA and EGI-InSPIRE, and one infrastructure MoU between the Universidade Federal do Rio de Janeiro (on behalf of the GISELA partners) and EGI.eu.

With these agreements GISELA becomes an EGI Infrastructure Provider, entitled to a seat on the EGI Operations Management Board. "This is an opportunity to expose Latin American needs to middleware development," says Marechal, adding that "a tighter collaboration with EGI translates into better gLite knowledge dissemination and access to useful tools."

GISELA brings together 19 beneficiaries (and 12 other parties) from 15 countries in Europe and Latin America for two-year project that 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 EGI.

> To provide Virtual Research Communities (VRCs) 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.

In addition, GISELA aims to provide support to both small user communities based at research institutions or working in small collaborations, and large Virtual Research Communities spanning Latin America and Europe. This includes providing training and tutorials to new users, operational support, porting new applications to the grid or managing digital identities and grid certificates.

GISELA's support services are provided by the project's team and in collaboration with other initiatives.

At its start, the GISELA project supports research communities from 11 scientific domains including bioinformatics, life sciences, high energy physics, civil protection, earth sciences, engineering, chemistry, e-learning, computer sciences, mathematics and fusion. But they are keen to point out that neither the number of user communities nor the number of scientific domains is restricted and that new communities and new applications are more than welcome in the project. •

Bernard Marechal, GISELA's project coordinator, signed two agreement with EGI.eu's Steven Newhouse (left) and an MoU with the e-ScienceTalk project represented by Catherine Gater (right).

More Information

GISELA project
http://www.gisela-grid.eu/

Driving EGI policy

Sy Holsinger reports on the outcomes of the policy workshop at the User Forum

credit: freefoto.com

What drives policy? Who defines it? What does it even mean? The EGI ecosystem is a complex web of actors ranging from resource, service and technology providers to a diverse user community covering a wide range of activities spread across numerous countries and regions. Whether on a European or National level, policies provide the clarity and the processes for interaction within the EGI ecosystem.

The EGI.eu Policy Development Team (PDT) supports the coordination of these policies for the benefit of all, whether they be 'technical' through open standards, 'operational' for operational policies and procedures, or 'strategic' by providing analysis and guidance around key topics of discussion.

It has been an exciting first year with the establishment of EGI.eu and the start of the EGI-InSPIRE project. The PDT has supported the establishment of ten policy groups, including a policy development process to support them, and has prepared key strategic reports around sustainability plans, integration of clouds and virtualisation into EGI, EGI's role within the Europe 2020 strategy, the suitability of the ERIC legal framework for EGI.eu and a standards roadmap, just to name a few.

As the first year of EGI was devoted to establishing the general policy framework, the next step forward and priorities of the coming year are geared towards reinforcing the connection with the National Grid Initiatives (NGIs) and European Intergovernmental Research Organisations (EIROs) participating in EGI. This will facilitate EGI-to-NGI and NGI-to-NGI policy development interactions and an exchange of best practices.

This context served as the basis for the first of a series of EGI Policy Workshops, organised at the EGI User Forum in Vilnius on 13 April 2011. The session focused on the results of two recent surveys conducted by the PDT together with the NGIs and EIROs. The goal of the first survey was to engage with the community to kick-start a flow of ideas on key policy issues, obtain updates on NGI legal and organisational status and understand the most pressing issues around sustainability moving into the future. The second survey focused on understanding the viewpoints about the suitability of the ERIC legal framework for EGI.eu, and presenting the preliminary results from the NGIs and EIROs themselves for which road should be travelled.

Overall, the participants' feedback on both the surveys and session proved extremely positive. The opportunity to showcase a number of achievements through the EGI policy activities performed at both national and international levels underlined the importance of these activities and why it is essential for their involvement in this area. Consequently, channels have been established in order to continue directly engaging the ecosystem to ensure that EGI *collectively* evolves in the right direction. •

For more information regarding the survey analysis, responses to NGI/EIRO main points, and a list of references to EGI policy documents:

http://go.egi.eu/NGI-SAR1

Operations architecture and latest infrastructure figures

In January 2011, the Operations Management Board (OMB) met in Amsterdam to revise and approve the operations architecture, based on the notion of **Resource Infrastructure**.
Peter Solagna and Tiziana Ferrari walk us through the reviewed operations architecture and unveil the latest figures on the production infrastructure

EGI provides access to **resources**, i.e. logical and/or distributed entities to be shared by end-users – for example CPUs, data storage, instruments and digital libraries. Resources are contributed by **Resource Centres**, also known colloquially as *sites*. The Resource Centre is the smallest localised or geographically distributed administration domain, where EGI resources are managed and operated. Resource Centres also provide the grid functional capabilities necessary to make their resources seamlessly accessible to authenticated and authorised users through common interfaces.

Resource Centres are interconnected by the National Research and Education Networks (NRENs) and GÉANT, and are usually federated to constitute a **Resource Infrastructure**. The legal entity responsible for its operation and integration into EGI is the **Resource Infrastructure Provider**.

The services needed to guarantee day-to-day operations are delivered by an **Operations Centre**. Locally, Operations Centres are responsible for supporting their Resource Centres, monitoring their performance, collecting requirements and for representing them in EGI's operations boards. Globally, the Operations Centre contributes to the development of the EGI operations roadmap and the evolution of EGI operations on behalf of its Resource Infrastructure Provider.

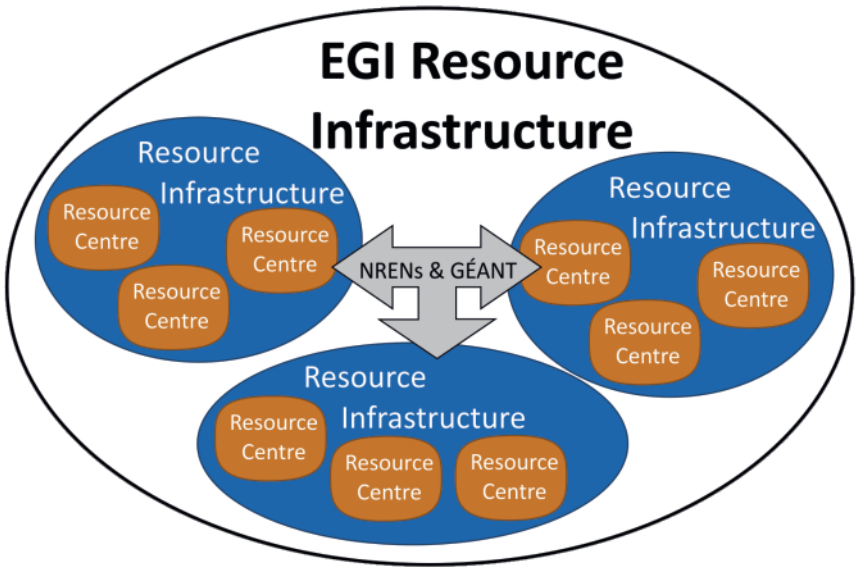

Local Services are complemented by EGI.eu Global Services such as the EGI Helpdesk, support, central monitoring and accounting and infrastructure oversight.

The sum of Local and Global Services in operations constitutes the **EGI Service Infrastructure**.

Resource Centres

At the end of March 2011 the EGI production Infrastructure included **332 Resource Centres** from 58 countries and CERN. Thirteen of these resources are contributed through non-EGI-InSPIRE partners: China, Pakistan and New Zealand (Asia Pacific Federation); Austria and Estonia (NGI_NDGF Federation); Belgium (NGI_NL Federation); Argentina, Brazil and Venezuela (IGALC Federation); Canada and China (Canada Federation); Brazil, Chile, Colombia and Mexico (Latin America Federation).

In the last 12 months the number of countries increased from 48 to 58: a 21% increase driven by the integration of infrastructures in the Baltic and South East Europe regions. Albania and Moldova are planning to be integrated during the course of 2011.

More Information

EGI Operations Architecture
https://documents.egi.eu/secure/ShowDocument?docid=218

Resource Centres
(EGI-InSPIRE partners)

March 2011:	**332**
Target (2011):	300

Operations Centres

At the end of EGEE-III the Resource Infrastructure was operated by 14 Regional Operational Centres (ROCs): Asia Pacific, Canada, Central Europe, CERN, France, Germany/ Switzerland, IGALC, Italy, Latin America, Northern Europe, Russia, South Eastern Europe, South Western Europe, and United Kingdom/Ireland.

This scenario has changed considerably during the first year of EGI-InSPIRE as the largest ROCs (Central Europe and South East Europe) were decommissioned. The EGEE ROCs have consequently developed into a much larger group of smaller Operations Centres, which typically serve a single country.

EGI currently counts **32 Operations Centres**: 22 are managed at a national level, one by a European Intergovernmental Research Organisation (CERN) and four by non-European Operations Centres: Asia Pacific, Canada, IGALC, and Latin America.

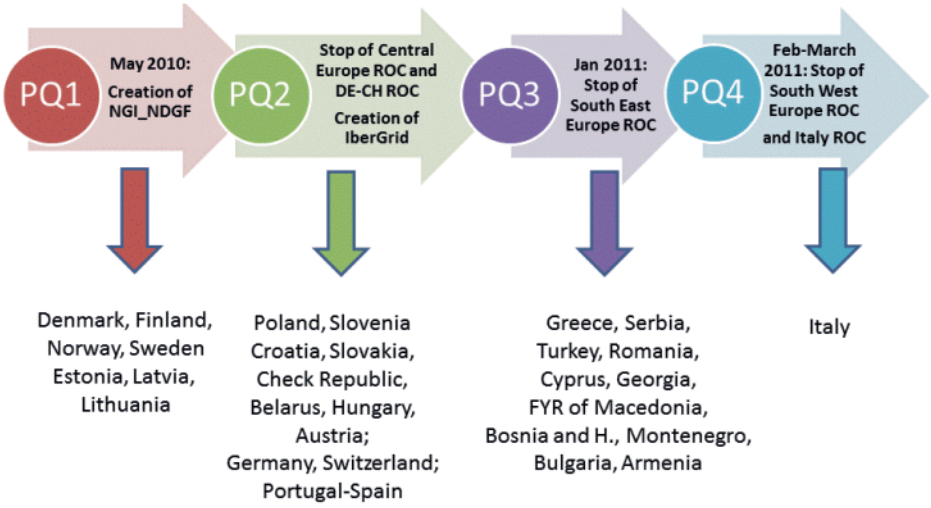

The break-up of EGEE's ROCs during the first four quarters (PQs) of the EGI-InSPIRE project.

Infrastructure capacity

The capacity of the EGI Resource Infrastructure is regularly monitored. The metrics used to make this assessment are the number of logical CPUs (cores) and the HEP-SPEC 06, which is a standard metric to assess the computing capacity of physical machines.

The increase registered from April 2010 – at the start of EGI-InSPIRE – amounts to 7.9 % for logical CPUs and 39.8% for HEP-SPEC 06. The total installed disk capacity reported is 101 PB while the total ins-talled tape capacity amounts to 80PB.

Logical CPUs

	May-July '10	Aug-Oct '10	Nov '10-Jan '11
EGI-InSPIRE partners	184,844	197,777	207,203
Target 2011: 200,000			

HEP-SPEC 06

	Capacity HEP-SPEC 06	Average HEP-SPEC 06/core
EGI-InSPIRE partners	1,873,698	9.1

HEP-SPEC 06 Reference:
http://go.egi.eu/hep-spec06

User statistics

At the end of March 2011, EGI had:

> **13,319 users** (9.5% increase from March 2010)
> **186 Virtual Organisations**, or VOs

High Energy Physics (HEP) accounts for 45.4% of the total number of users, followed by the 'infrastructure' VOs (a container of service or catch-all VOs) with 14.6%, multidisciplinary VOs (10.1%), Life Sciences (6.5%) and Computational Chemistry (3.7%).

Job averages

(Average number of million jobs submitted per month)

All VOs

May '09-Apr'10	**16.5**
May '10-Mar'11	**27.7**

Excluding HEP VOs

May '09-Apr'10	**0.554**
May '10-Mar'11	**0.909**
Target 2011	0.500

Reliability and Availability

The quality of operational services delivered by Resource Centres, Resource Infrastructures and EGI is measured with Availability and Reliability metrics, computed from the results of periodic tests performed at all certified centres.

Reliability

May'10-March'11	**91.9%**
Target 2011	90.0%

Availability

May'10-March'11	**90.7%**

Reliability and Availability stats:
https://wiki.egi.eu/wiki/Performance

Opinion: Finding out what researchers need

Simon Hettrick, from the Software Sustainability Institute, on a new strategy to gather user community requirements

(Photo courtesy of vitroid)

How's this for an idea: better research comes from investing money into what researchers need. It's the kind of brutally simple plan that would work – if only it was as easy to implement as it is to describe. Frequently, the difficulty lies not in giving researchers the help they need, but in finding out what they need.

The Software Sustainability Institute's (SSI) job is to help researchers with software. We might help develop better code, build communities around a researcher's software or simply make sure that other people are aware of the software and what it does. But before we can do our work, we need to know which researchers need our help and what they want us to do.

When it comes to requirements gathering, people frequently set about a comprehensive study. This means collecting and reviewing different (and frequently conflicting) needs from hundreds of people. A typical research community will have its theorists and experimentalists to consider and different perspectives from project leaders, researchers, managers and technicians. Surveys of this kind are prohibitively expensive, and can take so long to complete that the requirements have changed before it's finished.

The SSI is using a different approach: we've asked the researchers to come to us and tell us what they need. This will not lead to a comprehensive study of a community. Instead it will attract responses from a specific type of researcher: the ones who have problems to solve, are keen to try out new ways of working and are willing to work with people outside of their group. In other words, exactly the kind of people we should work with.

An example of this approach is the Collaborations Workshop, which we ran in March this year. The workshop brought together researchers, software developers and funding agencies, and got them talking about the problems they face in their work and their ideas for solutions. With such a diverse group, it would be almost impossible to set a fixed agenda. Instead, we asked people to suggest topics for discussion, these suggestions were voted on by everyone at the workshop and only the most popular topics made the agenda.

The focus of the Collaborations Workshop is on the changes that will help research. Each topic is discussed by a group, and it is the group's job to come up with a list of changes that will make life easier for researchers. The group's findings are then reported back to everyone at the workshop and made available online. This year, we generated a list of around 50 changes that would help researchers. These ranged from lobbying funders to make software sustainability a part of research proposals, to making software training resources more visible on the net.

By asking researchers to come to us with their suggestions, we end up talking to the kinds of people that work well with the institute. Focusing on this group means that we can efficiently create and prioritise a list of the work the institute should do, so that we can invest our resources in the projects where we can help the most. •

More Information

Software Sustainability Institute
www.software.ac.uk

The Collaborations Workshop
www.software.ac.uk/home/cw11

The SEERA-EI comparative analysis of national programmes

Dimitra Kotsokali, SEERA-EI dissemination manager

The South East European Research Area for e-Infrastructures – SEERA-EI project, compiled an in-depth comparative analysis of national programmes, and national and regional e-Infrastructures, to identify gaps in development and regional collaboration areas. The analysis is part of the SEERA-EI strategy to overcome obstacles and to identify policies related to e-Infrastructures, focusing on both short- and long-term actions that will provide a competitive scientific milieu for the region and Europe. The outputs will be used to formulate a collection of national best practices.

The document presents the framework of e-Infrastructures in terms of its main constituents (e.g. network, grid, HPC, middleware, cloud computing and green IT) and reviews pan-European (e.g. EGI, GÉANT, TERENA) and regional e-Infrastructure collaborations (e.g. the SEE-GRID project series).

The technological priorities and potential areas for future regional collaboration identified were:

> the stability and sustainability of network operations, specifically dark fibres;

> the emerging computing paradigms, such as HPC, cloud computing and Green IT;

> transnational access and regional end-user applications;

> ESFRI integrations;

> softer actions focused on training and human capital enforcement.

The study also highlights the importance of initiatives such as EGI, PRACE, DEISA or GÉANT to encourage different users in the deployment of e-Infrastructure services.
Participation of users, providers and developers in the development of e-Infrastructures is encouraged to obtain the most effective and fruitful services, which may trigger innovation.

The harmonisation of e-services and policies to improve collaboration and the creation of structured and long-term mechanisms to support e-Infrastructures have the potential to create superior economies of scale. These can provide the seamless e-Infrastructures required to produce high-quality research and training, to deploy technology transfer and best practices, to boost innovation, and to provide sustainability and global competitiveness to European e-Infrastructures.

To this respect, SEERA-EI has established liaisons with EGI, ESFRI, e-IRG, PRACE, WINS-ICT, e-InfraNet, and other initiatives, that have supported the sharing of the know-how in terms of policies, programmes and vision within the SEE region and beyond. By regular exchange of information with the main European actors in e-Infrastructure the region will be able to reduce the fragmentation and increase the coherence in the policies and strategies regarding e-Infrastructure programmes. •

More Information
www.seera-ei.eu

EDGI introduces new tool to submit 10,000 jobs to the grid

The European Desktop Grid Initiative (EDGI) project successfully demonstrated the use of the new MetaJob feature at the EGI User Forum in Vilnius. MetaJob enables users to submit jobs to the EDGI Desktop Grid through gLite using one single command and can be used to wrap or describe thousands of jobs into one single MetaJob.

"Desktop Grids are especially useful for applications like parameter sweeps where you have thousands and thousands of similar jobs," explains Ad Emmen, from AlmereGrid. "MetaJob is a command line 'script' that helps an EGI user to create these jobs for his/her application, and submit them automatically through standard gLite interfaces."

The efficiency of MetaJob was demonstrated with great success at the EDGI booth. During the day, the EDGI team submitted 10,000 jobs several times to show the EGI User Forum visitors how this can easily be processed by BOINC, without disrupting the gLite infrastructure.

For Jozsef Kovacs, from the MTA SZTAKI (Computer and Automation Research Institute, Hungarian Academy of Sciences), MetaJob makes it easier to handle a high number of jobs. MetaJob also allows "to increase the scalability of the service grid to desktop grid bridging services deployed by the EDGI project by decreasing the load of the service grid components – for example gLite wms," he adds.

From a technical point of view the solution is based on the well-known Service Cluster Grid to Desktop Grid bridge, where an extra MetaJob description file (containing the definition of a huge number of jobs) is attached to the job as input file. This MetaJob input file is recognised by the 3GBridge component, which then creates the jobs and inserts them into BOINC. This solution also works with XtremWeb on the Desktop Grid side or ARC on the Service/Cluster Grid side without any modification.

"This successful EDGI demonstration clearly shows that Desktop Grids can be used effortlessly for scientific day-to-day calculations," said Peter Kacsuk from the EDGI project. "The relevance of Desktop Grids for e-Science has been proven once and for all." •

More Information
http://edgi-project.eu/

Upcoming events

Technical Forum 2011 - Call for participation opens

The EGI Technical Forum 2011 will be held in Lyon, France between 19-23 September 2011, in colocation with Grid 2011 and the Open Grid Forum's next major European meeting - OGF33.

The programme committee is now calling for contributions from the European e-Infrastructure community, and their international collaborators, on a wide range of topics including (but are not limited to):

> EGI operations
> Technology used within EGI
> The support, work and tools of the EGI User Communities
> EGI policies
> EGI dissemination
> Collaborating projects

Abstracts should be submitted through the Technical Forum's Indico page until 3 June 2011.

The event will also feature an exhibition area open for organisations active within the EGI community (projects, NGIs, companies) to present their work. In addition, end-users, application and tool developers, operations staff and others are encouraged to participate by submitting abstracts for:

> Sessions
> Posters
> Demonstrations
> Workshops

The programme committee will review the contribution requests and allocate space accordingly, which may require different contributions to be merged.

Posters will be on display all week during the event and the selected demonstrations will have allocated time on the booths provided. The exhibition area will be open for organisations active within the EGI community to present their work.

Requests for co-located workshops will need to identify how the workshop organisers will ensure a high-quality programme that is relevant to the EGI community. •

Call for participation is open

Deadline for abstracts: 3 June 2011

More Information

http://tf2011.egi.eu/

Gearing up for IBERGRID'2011 in Santander

The 5th Iberian Grid Infrastructure conference – IBERGRID'2011 – will be held in Santander, Spain between 8-10 June 2011, at the Palacio de la Magdalena.

The event, hosted by the Advanced Computing and e-science Department at the Institute of Physics of Cantabria (CSIC-UC) and the Polytechnic University of Valencia, will focus on the ongoing grid projects in development by Portuguese, Spanish and Latin American institutions.

"The IBERGRID conference has been established as the meeting point where researchers, technicians and developers of the area of distributed computing from Spain and Portugal gather once per year," says Isabel Campos, chair of the conference's organising committee.

"The attendance is very similar to that of a user forum – about half of the presentations are devoted to scientific applications running on the EGI infrastructure," she says adding that software developers and site administrators also come to IBERGRID, which has become a sort of a tradition for the Iberian community.

The IBERGRID'2011 programme is varied and covers topics such as green IT, large data repositories, LHC Computing Grid and Tier-0 topology, interactive supercomputing, applications for high performance networks, ESFRI implementations in the Iberian area or applications for volunteer computing.

This year the event will also feature special tracks on virtualisation techniques and its applications, as well as on advanced management of computer centres. •

More Information

www.ibergrid.eu/2011/