




European Grid Infrastructure

Inspired

Winter 2012

News from the EGI community



EGI.eu's second anniversary // page 1

CF 2012: Programme preview // page 2

CF 2012: Training workshops - go with the (work)flow // page 3

Case study: Hunting for new viruses // page 5

Music research and grid computing // page 6

Sustainability: the next steps // page 7

European Research Area – No investment without reform // page 8

A positive outlook for availability in EGI // page 9

EGI going social // page 10

Mascot found! // page 10

This Issue

Welcome to the Inspired Winter 2012 issue!

This time we are getting ready for the upcoming Community Forum in Garching, Munich, with a preview of what's interesting by Richard McLennan and an overview of the Training Workshops by John Walsh. But that's not all:

- > Steven Newhouse celebrates EGI.eu's second anniversary
- > Steve Brewer has a look at what is being done in music computing research
- > Tiziana Ferrari has good news about availability in EGI
- > Sy Holsinger summarises the conclusions of the latest Sustainability Workshop
- > Damir Marinovic was our eyes and ears at the European Research Area conference in Brussels
- > Neasan O'Neill tells us how EGI is engaging with several social networks
- > Viviane Li announces who won the mascot competition
- > and I report on how grid computing is helping the search for new virus species.

As always, if you want to contribute ideas, suggestions or stories to the newsletter don't forget to let me know!

Sara Coelho
sara.coelho@egi.eu



The cold spell in late January / early February brought misery to those hailing from warmer climates, but our Dutch friends were very happy to go skating on the frozen canals. As you can see from this 19th century painting by Andreas Schelfhout, ice skating is a hobby with a long tradition.

EGI.eu's second anniversary

Steven Newhouse looks back on a year of achievements

On 8th February, EGI.eu celebrated its second anniversary as a not-for-profit foundation dedicated to linking researchers with the computing resources they need for their work – happy birthday!

Our second year of activity was full of productive discussions, important milestones and significant achievements. Here are a few highlights:

> We have signed **18 MoUs** with partner projects, resource infrastructure providers and virtual research communities.

> We organised two successful **flagship events**: the User Forum in Vilnius and the Technical Forum in Lyon. With more than 655 participants, the Technical Forum now holds the record for the biggest European grid event to date.

> The EGI-InSPIRE and e-ScienceTalk projects, both coordinated by EGI.eu, passed their **first reviews** by the European Commission.

> We have recognised a new role in the community – the **NGI International Liaison (NIL)**. The NILs will bring our community closer together and will act as main contact point between EGI.eu, each NGI, and other NGIs within EGI.

> We developed the idea of the **Virtual Team** – essentially, a framework for 'getting things done'. The Virtual Teams, enabled by the

NIL coordination role, will allow us all to focus on the issues that are important to the community.

At two years old, EGI.eu is still a young organisation, but it's never too early to start thinking about the future. Our third year will focus on our vision for the future as a global organisation and how we can make sure that the EGI ecosystem will continue to provide sustainable computing services to scientists and researchers all over the world for many years to come. •



CF 2012: Programme preview

Richard McLennan tells us what's in store for the Munich event



As you read this, the EGI Community Forum will be nearly upon us and if you have not already registered to attend, now will be a good time to scribble yourself a reminder on a 'Post-It' and quickly book yourself in.

The EGI forums are well established as key fixtures in the grid community's calendar, with consistently growing attendance that now surpasses 600 delegates. This forthcoming event can be expected to be equally popular and is being run from 26 to 30 March at the Leibniz Supercomputing Centre in Garching near Munich, together with EMI who are holding their second Technical Conference. The EGI Community Forum will be preceded by the European Globus Community Forum on Monday 26 March at the same venue.

Situated on the northern outskirts of Munich and in the economic centre of southern Germany, the Leibniz Supercomputing Centre serves Munich's universities and the Bavarian Academy of Sciences and Humanities. Perhaps more significantly, from the EGI perspective, it provides high-end computing facilities for the scientific community in Germany and beyond.

Nearly 170 submissions have been accepted into a technical programme that focuses at its core on the sustainability of the coalition of resources that make up 'the grid'. More than just an opportunity to share knowledge, the event will be an important milestone for the

research community to state its needs and its vision for the future of our e-infrastructure and for EGI.eu as its evolving coordinating body.

The Forum has been arranged into five broad tracks covering Users and Communities, Software Services for Users & Communities, Middleware Services, Operational Services & Infrastructure and lastly, Coordination & Communication. Submissions cover a very wide range of new material and provide ample evidence of the important role that the grid now plays across the whole field of research.

Attendees will hear about the challenges and successes of running DNA sequencing experiments on the grid and how e-BioInfra is used daily for analysing high throughput sequence data on the Dutch grid. The grid enables bioinformaticians and biomedical users to scale up their analysis and be flexible in their choice of analysis tools to use in a scientific workflow.

We'll have a report on Workflow and Data Management for Nuclear Magnetic Resonance with a focus on users who are not specialists in computational structural biology and thus need simple methods to obtain reliable results. As an extreme example, a demonstration using data sonification (algorithmic composition) as an artistic tool for musicians will serve to show the diversity of use that can still be expected from distributed computing. •

Confirmed keynote speakers

- > **Steven Newhouse**, EGI.eu
- > **Kostas Glinos**, GÉANT
- > **Alberto Di Meglio**, EMI
- > **Wouter Los**, ENVRI project
- > **Daan Broeder**, DASISH project
- > **Alexandre Bonvin**, WeNMR
- > and **Horst Zuse** who will give a special talk on Friday on the history and the origin of the first computer

<http://cf2012.egi.eu>

More Information

The EGI Community Forum 2012 and the 2nd EMI Technical Conference will take place 26–30 March

Programme:
<http://go.egi.eu/cf12>

Registration:
<http://cf2012.egi.eu/registration/>

CF 2012: Training workshops - go with the (work)flow

John Walsh introduces the community training workshops

The forum will see experts across the scientific domains gather to both publicise new results from their grid-based computational models, and to display the wealth of tools available to help experienced grid users and unversed scientists to exploit the EGI distributed computing infrastructure (DCI) with greater ease.

The Heavy User Communities (HUC) training sessions and workshops shall be as diverse as ever. In addition to a range of Tools and Shared Services workshops, this year will also see a greater focus on Workflows and Portals.

Here is a preview of the up-and-coming workshops.



The workshops will take place in Seminar rooms 1 and 2 at the Faculty of Mathematics and Informatics - slides provided for easy access.

Workflow Management Workshops

Workflows describe a set of linked tasks or processes that yield an outcome. Scientific workflow management systems are particularly good at handling a series of complex tasks using large datasets in numerous data formats. A workflow can be described once, but used many times with different input parameters. Portals, on the other hand, can provide user-friendly, web-based access to workflow management and job submission systems. Indeed, Workflow Management Systems and Portals, play a huge role in managing scientific workloads across all the scientific disciplines, and are expected to have a greater role to play in our daily lives as more and more people move to tablet-based computers.

Kepler

The Kepler Scientific Workflow Tutorial will be 'hands on', with attendees being introduced to: basic workflow concepts (e.g. frameworks, actors, directors); how to build basic workflows; how to use relations,

paths and synchronization, if-else, loops; and finally, how to use advanced features, such as grid actors, in simple and complex workflow scenarios.

WS-PGRADE

The two-part tutorial on WS-PGRADE/gUSE presents an open-source, web-based workflow management tool that enables users to exploit several grid systems, use multiple grid technologies, and to access other DCI infrastructures. It has a feature-rich and easy-to-use interface. The first session focuses on User and Application development training and will cover: WS-PGRADE workflow concepts; parameter study features and support in workflows; the WS-PGRADE portlets; access to other DCI infrastructures; scientific application areas using WS-PGRADE. The second session will be at a more advanced level and covers WS-PGRADE/gUSE portal development. It is primarily aimed at scientific communities, NGIs or portal developers who would like to create customised WS-PGRADE based

science gateway applications for their users. This tutorial will introduce the attendees to the Application Specific Module (ASM), a flexible API layer to assist the creation of customised science gateways.

SHIWA

This tutorial introduces SHIWA (Sharing Interoperable Workflows for Large-Scale Scientific Simulations on Available DCIs), a tool aimed at handling a wide variety of scientific workflows in a platform neutral way.

SHIWA allows workflows described in one format to be translated into other formats, and enables workflows to run on several DCIs. SHIWA also provides a searchable workflow repository, thus allowing e-scientists to reuse them in their own workflow applications.

This is important because communities that have developed workflows may be susceptible to vendor 'lock-in' – i.e., prevented from using them in other workflow management systems, and tied into particular computing infrastructures.

Tools and Shared Services Tutorials

The Tools and Shared Services tutorials look at a range of tools that can be exploited by all grid users.

CERN VM FS (CVMFS)

The CERN Virtual Machine File System (CVMFS) is a caching, http-based, read-only, scalable file system optimised for delivering experiment software to (virtual) machines. Other primary features include: standards-based technologies; file system security (SHA1 check via HTTPS); file-based duplication; pull-based solution for deploying experiment software at each site.

This session is aimed at user communities wanting to explore the use of CVMFS as a new, efficient and flexible means to manage and deploy their experiment software on the grid. In particular, the session will

discuss the server-side configuration, and should be of interest to resource centre administrators and user communities alike.

Parallel Computing Workshop

This workshop (which is included in the 'Making DCIs work for you' session) will concentrate on the current state of the art for successful parallel job submission on the infrastructure. The topics covered will include: description and submission of parallel jobs; techniques for successful job submission, and pitfalls to avoid; execution of MPI applications using gLite, ARC and Unicore middleware stacks; support for other parallel application frameworks (OpenMP and hybrid MPI/OpenMP); use of GPGPUs.

StratusLab

With over 50 attendees, the StratusLab tutorial was one of the success stories at the EGI User Forum 2011. Participants will learn about cloud technologies in general, and discuss the distinction between Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS) Infrastructures. The workshop will discuss how infrastructures based on the StratusLab distribution can be integrated with EGI and how cloud services complement grid services. Practical exercises will teach the participants how to launch virtual machines, customise their computing environment, share those environments with others, manage virtual disks, and define complete services. •

| | Tuesday | Wednesday | Thursday | Friday |
|-------------|---------------------------|-------------------|---|---------------------------------|
| 11:00-12:30 | Kepler tutorial SR2 | Cern VM FS SR1 | StratusLab workshop SR1 | SHIWA workshop (part 1) LRZ2 |
| 14:00-15:30 | WS-PGRADE (part 1) SR2 | | | SHIWA workshop (part 2) LRZ2 |
| 16:00-17:30 | WS-PGRADE (part 2) SR2 | | Parallel computing workshop (Making DCIs work for you,pt2) SR3 | |

SR1, SR2 and SR3 = Seminar Rooms 1, 2 and 3. Please note that timetable and room allocation may be changed - check <http://go.egi.eu/cf12> for updated information.

Case study: Hunting for new viruses

Sara Coelho finds out how grid computing helps scientists to discover new diseases

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

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

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

Looking for viruses

Lia van der Hoek and colleagues from the Virus Discovery Unit, at the Academic Medical Centre of the University of Amsterdam (AMC), has been working on VIDISCA – a method to spot new viruses from previously unidentified genetic sequences.

The hunt for new viruses starts at the hospitals, with the collection of nose and throat swabs from patients. Back at the lab, the first step of the VIDISCA method is to remove residual cells and other biologic material, to enrich the sample's viral genetic material. The genetic sequences in the sample are then amplified with standard techniques.

Success first came in 2004, when the team reported the discovery of the coronavirus NL63, which is implicated in croup – a disease that causes throat swelling and coughing in children under six-years old. The virus was spotted in samples taken from a 7-month old baby, admitted to a Dutch hospital with symptoms of acute respiratory infection.

The genome of the coronavirus NL63 was sequenced and the analysis, published in *Nature Medicine*, showed that the virus was a new species with distinctive features.

The VIDISCA haystack

Over the past six years Lia's work on the VIDISCA method has benefitted from the so-called next generation sequencing techniques. The improved version, dubbed VIDISCA-454, was introduced in 2009.

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

The National Centre for Biotechnology Information (NCBI) has created the BLAST tool to compare given sequences with databases via the web. But uploading data from VIDISCA-454 to this portal proved to be virtually impossible, given that the average experiment produces approximately 400,000 sequences.

Grid computing to the rescue


Looking for a solution, Lia contacted Antoine van Kampen, head of the AMC bioinformatics department,

who assigned Barbera van Schaik to the problem. Barbera developed a workflow – the sequence of computational steps required to perform an analysis – to allow BLAST to run on grid computing resources of the Dutch e-science grid.

The workflows and databases were made available to the Virus Discovery Unit via the e-BioInfra platform developed and operated by the e-bioscience group of Silvia Olabarriaga.

With these tools at hand, Lia's team is able to analyse a VIDISCA-454 experiment within 24 hours compared to weeks of intensive manual work. A test with 1444 samples produced 4,783,684 sequences and showed that the analysis can be repeated within 14 hours, compared to 17 days if it were to run sequentially on a local server.

Research is now ongoing on the meaning of the 4,000,000-plus sequences identified by VIDISCA-454. Lia now hopes to find new viruses linked to more respiratory infections and other serious diseases. •



Respiratory infections such as croup (that can cause a narrowed trachea, shown by arrow) are a major child health problem, but many are caused by still unknown viruses. (Illustration: wikicommons / Frank Gaillard)

Music research and grid computing

Steve Brewer investigates how grid computing is contributing to musical research

The music computing research community is fast following in the footsteps of physicists, astronomers and life-scientists in their use of the grid to radically extend the scope of their work. Here are three recent examples of how this emerging research field can benefit from what the distributed power of grid computing has to offer.

ASTRA

The ASTRA project (Ancient instruments Sound/Timbre Reconstruction Application) builds on existing work in a number of diverse fields. Archeological remnants and documents have been used to reconstruct lost ancient instruments, while modelling a sound synthesis techniques recreated the acoustics of these devices. But it was the addition of grid computing technology that allowed the researchers from Parma and Salerno, led by Domenico Vicinanza, to run computing-intensive digital signal processing algorithms in vastly reduced times. The result was that instruments such as the barbiton and the harp-like epigonion could be heard again and then, more recently, featured in original performances.

PLOrk

The curious-sounding acronym PLOrk stands for the Princeton Laptop Orchestra. PLOrk is not the only laptop orchestra, but the Princeton ensemble, originally established by Dan Trueman with fellow professor Perry Cook, continues to set the pace. As with traditional orchestras, the beauty of the digital version is that humans can and must continue to contribute significantly to the timbre, dynamics and rhythm of the performance. But as with a fully analogue orchestra there are many other tasks involved too, such as distributing the appropriate

musical scores to the right musicians at the right time. The full complexity of managing this distribution of material in a timely manner to the appropriate devices – laptops, Wii-motes, iPads and the like – required a new approach. Stephen Beck and colleagues behind the Grid Enabled Deployment for Laptop orchestras (GRENDL) project turned to the well-established SAGA grid framework to coordinate the participants in the laptop orchestra. This has taken what was a quirky but growing niche into an unbounded opportunity for musical growth thanks to the flexibility and scalable benefits of grid computing.

Peachnote

Peachnote is a relatively new project based on the work of mathematician and conductor Vladimir Viro, who is now studying Computer Science at the Ludwig-Maximilians University in Munich. Thanks to the support of a Google Research Award, Vladimir and his colleagues have developed a search engine that does for music scores what Google Books does for the corpus of digitised literature on the net. So now, if you have a phrase or motif that you want to investigate you can enter this into their system which rapidly searches through a library of 160,000 musical scores. This clever piece of software engineering brings fascinating questions to life: where do key phrases in important works come from? Who influences certain composers and how do phrases and melodies travel through time and space? The service is online, freely available and has already been accessed by tens of thousands of inquisitive musical minds, both professional and amateur. The computational power is currently delivered through the Apache Hadoop framework coupled with a



Music to grid ears - study of a cello.
(Illustration: wikicommons / MichaelMaggs)

user-friendly Google-technology based user interface.

Future?

Although these three projects have all started to gain publicity due to their individual originality, they just scratch on the surface of what could be achieved with the power and flexibility of the grid to scale the work to greater levels. What is also interesting about these endeavours is that, although unconnected, they collectively represent the foundations for a significant new community for EGI to work with. Together, we can create a set of grid-based resources that others can build upon to do further ground breaking-research across the music computing spectrum. •

More Information

ASTRA project
<http://www.astraproject.org/>

PLOrk
<http://plork.cs.princeton.edu/>

Peachnote
<http://www.peachnote.com>

Sustainability: the next steps

Sy Holsinger reports on the outcomes of the Sustainability Workshop in Amsterdam

Since it first appeared in the 1980s, the use of term 'sustainability' has grown exponentially; however, it has been difficult to pin down what it means and what people understand by it. Within the context of EGI, sustainability is seen as 'Establishing an enduring, open ecosystem with interacting but independent components which can fail and be replaced by others without damaging the ecosystem as a whole.' This ensures the availability of EGI for decades to come.

Why sustainability?

The European Commission appears no longer be prioritising funding for operations and maintenance of the grid infrastructure. Middleware development projects, such as EMI and IGE, are ending in Spring 2013 with no immediate plans for any follow on projects. Groups dependent on such middleware services - the infrastructure providers and individual research communities - are exploring what the support options might be. EGI's move towards a federated cloud infrastructure will change both the dynamic of the services that can be deployed by a research community to support their users and the services EGI needs to provide to facilitate this. In practice, this means that continued support for the services consumed by a researcher will increasingly depend on the commitment of that community to support those individual services in a sustainable way.

Last month, we welcomed more than 60 representatives from the EGI community for a three-day workshop (24-26 January) to continue discussions about the sustainability of the European Grid Infrastructure.



This workshop built on the first meeting held last September at the EGI Technical Forum in Lyon, where a variety of high-level topics were covered, including the introduction of business models for e-Infrastructures.

The January workshop focused more on in-depth discussions around vision and strategy, research community needs, and service development. Overall, the workshop allowed the research communities to identify key services they expect from the infrastructure providers and how they need them to be sustained, for example monitoring, accounting, security, web access to resources, marketing and outreach, and collaboration services.

The discussion on the sustainability of the EGI ecosystem as a whole was especially productive, which overall:

- > Established a major communication channel and a forum for the community to discuss the future evolution of EGI.

- > Streamlined and refined EGI's 'core services' identifying which services are community specific and which are generic to all, and those that are user facing to those that are infrastructure-facing.

- > Progressed the discussion of the service offer and added value of EGI and the NGIs. This will help to clarify the positioning of EGI versus organisations such as PRACE or

Amazon and enable the development of a better marketing and outreach strategy.

- > Identified the need for proactive EGI and community strategies to improve the relationship and engagement with ESFRI projects.

- > Explored resource allocation models for research communities new to EGI in the various NGIs and discussed potential methods for engaging them with trial resources and ways of ensuring a smoother transition between initial contact and infrastructure usage.

- > Understood the importance for application integration for research communities new to EGI, that sometimes lack the technical skills required to get applications up and running and moved towards clarifying who is responsible and/or able to provide it. •

More Information

The workshop webpage includes summary slides for each breakout session:
<https://www.egi.eu/indico/event/709>.

All information gathered through the workshop will be fed into the EGI Strategic Plan currently in production with the next major discussion to take place at the EGI Community Forum on Tuesday 27 March: <http://go.egi.eu/cf12>

European Research Area – no investment without reform

Damir Marinovic was at the ERA conference in Brussels

Not even a general strike halting all public transportation in Brussels and extreme weather conditions could stop 400 enthusiastic participants from joining one of the year's major events for European science, research and innovation held on 30 January 2012. I was among the lucky ones who managed to get to the conference venue.

The European Research Area (ERA) Conference 2012 – Fostering Efficiency, Excellence and Growth – was an opportunity to gather together major stakeholders, present preliminary results on European Commission public consultation and engage in policy debates.

The conference was a result of a consultation process launched in the end of 2011 to gather feedback on how to identify and overcome the gaps and bottlenecks on the road towards a single market for research and innovation by 2014. The key issues discussed were related to researchers' careers and mobility, research infrastructures (RIs), knowledge transfer circulation, open access and the ERA partnership.

EGI.eu, on behalf of the EGI, submitted a position paper on the ERA Framework that contained 15 key recommendations for ERA success.

Overall, the public response was overwhelmingly supportive for developing the ERA, with some interesting statistics - 70% of respondents considering RIs as a very important area, and 81% were in favour of increasing public funding and the role of the EU. Respondents also favoured higher participation of stakeholders in the ERA to ensure acceptance and ownership of policies agreed at the EU level.

The conference was opened by the European commissioner for

Research, Innovation and Science, Máire Geoghegan-Quinn. In a quite thought-provoking speech, the commissioner stressed that the EU is losing the battle with the USA, China and other emerging economies in an increasingly competitive research landscape. This poor track record of achievement in world-class research, coupled with zero economic growth, means that European research cannot sustain a "business as usual" mentality. Only comprehensive reform can bring investments and results. Ms Geoghegan-Quinn stressed the need to avoid duplication of infrastructure investment at a national level and emphasised the broader role of a new ERA-partnership, with a well-defined roadmap, based on common objectives and realistic deliverables.

The Research Infrastructure session led by Beatrix Vierkorn-Rudolph, the ESFRI chair, was of great interest for RIs and e-Infrastructures. It was also the perfect opportunity to provide participants with a copy of EGI position paper.

From an EC perspective, RIs need to demonstrate clear added value for innovation in Europe. The new EU legal framework, the European Research Infrastructure Consortium (ERIC), should also lead to clarification of internet protocols and open access rules. Speakers agreed that it is necessary for the ERA to "go digital" to exploit the full potential of e-Infrastructures. I pointed out the EGI recommendation to build an ERA Knowledge base including all European RIs, national and European research groups in order to help the research communities to build collaboration with e-Infrastructures.

Gudmund Host, the e-IRG chair highlighted the lack of attention to e-



The European Research Area will spark a new era of innovation in Europe.
(Illustration: wikicommons / Sebastian Ritter)

Infrastructures, especially overlooked in the preparation of new RIs. Only through e-Infrastructure development can research communities be efficiently connected with industry. During the discussion, I underlined one of the EGI recommendations that calls for the EC to address the legislative uncertainties that hinder commercial use of RIs and e-Infrastructures.

To conclude, the conference was a success in terms of fruitful discussion, networking and promotion of the EGI position. The next steps will be to submit a final EC proposal on priorities in the ERA Framework by June 2012. EGI is looking forward to establishing itself as a one of the major stakeholders in the creation of the online ERA. •

More Information

EGI position paper on the ERA Framework:
http://go.egi.eu/EGI_ERA

A positive outlook for availability in EGI

Tiziana Ferrari collects the latest figures

The quality of core grid services deployed by Resource Centres has been measured since 2008 through availability and reliability metrics. These are computed from tests performed periodically at all certified centres through the Service Availability Monitoring framework (SAM). These metrics express the level of functionality delivered by grid services and are used to flag up areas to improve in the infrastructure.

Certified Resource Centres guarantee 70% availability and 75% reliability for their services - the minimum values accepted for a Resource Centre, as defined in Operational Level Agreements established with EGI.eu.

Increasing the overall performance delivered to users has been an ongoing effort since the introduction of service level management. Availability and reliability averaged per quarter across the infrastructure have been steadily increasing from 2008 by about 1% per year, reflecting the growing maturity of the infrastructure.

Even if performance is increasing, availability and reliability figures can

be affected by changes (e.g. upgrades in software, new production sites being certified). For example, the reduction recorded from August to November 2011 (Fig.1) reflects the transition of various large federated Operations Centres to a set of operationally independent NGIs. EGI currently includes 347 certified Resource Centres. This number has increased by 9.5% since May 2010. Given this continued expansion, there is an increasing need for a managed process to gradually and safely bring new Resource Centres from test to production. This is now part of the operations roadmap for 2012.

User experience, however, does not depend only on resource-access services, but also on other top-level grid services operated by NGIs/EIROs. For this reason, in September 2011 the performance measurement framework was extended to include core grid services operated by the NGIs. The set currently includes the information discovery services (top-BDII)s and will be gradually extended

to workload management systems, file catalogues and virtual organisations management services, for example.

Currently, 77% of the Resource infrastructure Providers operate top-BDII services with an average monthly availability that exceeds 97%. France, Greece, Italy, The Netherlands and Turkey delivered steady 100% availability from September 2011 to January 2012 (Fig.2).

Starting in January 2012, NGIs whose top-BDII service availability does not reach 99%, will be assisted to define a plan for service improvement. •

Availability of services (or a sites) is the percentage of time that the services were actually up and running $(\text{uptime} / \text{total time}) \times 100$.

Reliability is the percentage of time that the services (or sites) were supposed to be up and running, excluding scheduled downtime $(\text{uptime} / (\text{total time} - \text{scheduled time})) \times 100$.




Figure 1: Quarterly availability and reliability of resource centres averaged across EGI from May 2010 to January 2012.




Figure 2: Classification of NGIs by the monthly availability provided by their top-BDII services. Monthly availability is averaged over the period September 2011 – January 2012.

Data collection with thanks to Peter Solagna/EGI.eu and Giorgos Fergadis/AUTH.

EGI going social

Neasan O'Neill

Want to know what we're really up to? Keep an eye on Facebook and Twitter.

EGI now has new Facebook pages and Twitter feeds to allow the community to know what we are working on, but more importantly to let EGI talk with its community.

EGI has been using various social media channels for a while, but we have now improved our presence by giving a dedicated voice on Facebook and Twitter to three of the teams within EGI.eu:

- > Policy and strategic planning
- > User community support
- > Technology and operations

These teams can now engage directly with their community, complementing the general EGI

information that has always been provided through social media. Bringing together the news, information and discussions in each area means that everyone can get involved and understand where EGI is going. These new channels will allow the teams to report from meetings, events and discussions as they happen, so everyone can get involved and have their say about the issues that matter to them.

The most important part of this, however, is you the community. We need you to tweet, like, follow, post and engage with the teams. Tell us whom we should be following. What we should be taking about? Where should we be looking next? What do you want from EGI? •

To find all the information you need about the accounts go to

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

Mascot found!

Viviane Li

At the EGI Technical Forum held in Lyon in September 2011, EGI.eu announced a world-wide idea competition to find a mascot for the European Grid Infrastructure (EGI). Two prizes were available – a Grand Winner prize, awarded by a judging panel at EGI.eu and a People's Choice prize, awarded to the most popular entry voted by you via an online poll.

The competition attracted entries across Europe and beyond, from the United Kingdom to as far afield as India, submitted by contestants from within and outside the immediate grid community. Ideas ranged from the abstract to the tangible, from a flame to a characterful sheep (who even lived out his busy social life on EGI's Facebook page for a period of time).

In the end, it was Pegi the peafowl, submitted by Stevy Georgiadi from Greece, who won hearts and votes in the online poll, winning the People's Choice prize.

The Grand Winner was awarded to Monica Brondi from Italy. Monica's idea of an ant and accompanied description demonstrated the most complete concept amongst the candidates, with many qualities aligned with EGI's. The use of a societal animal belonging to a network and a community was particularly suited to the nature of distributed computing, which can

also mirror the nature of research work.

Prizes will be awarded to the winners at the EGI Community Forum held in Munich next month.

Congratulations to the winners and thanks to all contestants for their effort, creativity and help in our mascot search! •

