



European Grid Infrastructure

Inspired

Winter 2011

News from the EGI community



New features in EGI Application Database

// page 1

User Forum 2011: Programme preview

// page 2

NGI profile: Lithuania's LitGrid

// page 3

WeNMR: an infrastructure for NMR and structural biology

// page 4

Project profile: Initiative for Globus in Europe

// page 5

Gathering requirements

// page 6

Request Tracker dashboard

Collecting operational requirements

News in Brief

// page 7

Venus-C opens clouds for research

SEE-ROC closes door on operations

The EGI blog is looking for contributors

Upcoming events

// page 8

User Forum key dates

TERENA conference preview

This Issue

Welcome to the Winter edition of Inspired!

This issue brings you the profiles of

- > **LitGrid**, the Lithuanian NGI and co-host of the EGI User Forum
- > the **Initiative for Globus in Europe**, the first technical provider to sign an MoU with EGI, and
- > the **WeNMR** virtual research community.

But that is not all...

- > Viviane Li looks at the programme of the upcoming **User Forum**
- > Carrie Solomon previews the **TERENA** networking conference
- > Peter Solagna & Gergely Sipos update the **requests collection process**
- > William Karageorgos announces the latest **AppDB** release
- > and Neasan O'Neill reports on the **Venus-C** open call

If you want to contribute with ideas, suggestions or stories to the newsletter please let me know!

Sara Coelho

sara.coelho@egi.eu



It's cold in Amsterdam but we are not hibernating!

New major features in EGI Application Database

William Karageorgos from the AppDB development team reports

The EGI Applications Database (AppDB) stores tailor-made computing tools for scientists to use. The database embraces all scientific fields, from resources to simulate exotic excitation modes in physics, to applications for complex protein sequences analysis.

After the successful testing of the new version of AppDB with several NGIs, AppDB v1.1.0 was released on 4 February 2011.

The new release includes several functionality improvements and bug fixes and is available on the Applications Database website. New major features in this release are:

- > Improved and simplified write access: Every EGI SSO account holder can register applications and create developer profiles in the system.

- > Application Programming Interface (API)-based access: The system provides access for users of third party software through a new, REST web-API. The API allows read-access for external sites and it is currently in beta mode. Full write-mode access from programming interface is expected in May 2011.

Through the API, NGIs and VOs, for instance, may be able to provide their own localised, custom interfaces to the applications database service, which could, for example, display entries relevant to their country or discipline only, or in their own native language.

Data currently provided includes information on applications, developers, and related publications, and supports filtering via a number

of parameters such as discipline, country, virtual organisation or middleware.

Documentation about the new release and about upcoming releases and features is available within the system. Users can provide feedback about AppDB through GGUS, and are welcome to submit requirements for the tool through the EGI channels. •

More Information

<http://appdb.egi.eu>

User Forum 2011: Programme preview

Viviane Li finds out what the programme has to offer



What is your interest in the European grid? Whatever the answer, make some space in your diary from 11–14 April. This is when the first EGI User Forum takes place in Vilnius, Lithuania. No matter what your role is within the grid community, you will find sessions relevant to you.

“The idea this year is to combine the EGI [European Grid Infrastructure] user forum with the EMI [European Middleware Initiative] technical conference, bringing together the two aspects of infrastructure and middleware,” says Alberto Di Meglio, project director of EMI and one of the 29 programme committee members. “We have a balanced mix, so users can see how these aspects are integrated to provide a service directly usable for them.”

Even the best infrastructures and middleware stacks (software packages that link together the hardware resources) are incomplete without you – the users, developers and system administrators. With EGI.eu now acting as a point of coordination for the NGIs, the grid community can have a more formal voice in this year’s forum.

“There are many training sessions delivered by partners in the grid community as a part of the core programme,” adds Gergely Sipos of EGI.eu, another committee member. “The new programme is designed to empower different communities. This is enabled by the new EGI structure.”

EGI.eu will deliver sessions on ser-

vices for virtual research communities, highlight success stories, discuss dissemination and policy. There are different presentations, tutorials and workshops aimed at researchers, programmers, developers, system operators and technology providers. The range of topics and the variety of formats will allow wider participation than before. You can pick and mix sessions according to your interests.

EMI, one of the forum’s co-hosts, will run their first ‘meet the experts’ session. Di Meglio explains what they hope to achieve: “We want all users to have the opportunity to talk with us informally about their needs and formulate some concrete requirements for future development. We want to create a strong and more direct link between us – the technical providers – and users of the technology.”

EMI will also be launching ‘EMI-1’, a software designed to provide compatibility between middleware stacks – gLite (Lightweight Middleware for Grid Computing), ARC (Advanced Resource Connector), UNICORE (Uniform Interface for Computing Resources) and dCache.

With 154 abstracts accepted, there is a packed timetable of sessions delivered in parallel from six or seven conference rooms each day, not counting the poster and demonstration areas. The programme includes talks about the grid’s applications in biomechanical markers, molecular simulations, a paleobiology database,

electrophysiology, the Large Hadron Collider, and simulations for social science.

Coming to the forum will also help you to use the infrastructure in better ways. You can find out how to optimise workflow, implement data management technology, self-evaluate e-Infrastructure funded projects, attend user support services and a wide variety of workshops. You can also explore aspects of Distributed Computing Infrastructure, Desktop Grid, Virtualisation and Cloud computing.

The diversity of the forum is a great opportunity to network with new communities, as well as to catch up with members of your own Virtual Organisation face-to-face.

See you in Vilnius!•

For more information:
<http://uf2011.egi.eu>

Registration is open until 1 April;
early bird fees end 28 February.

The EGI User Forum 2011 is organised by EGI.eu, Vilnius University and LitGrid, with the EGI-InSPIRE and EMI projects and the support of local secretariat BAIP (Baltic Amadeus Infrastruktūros Paslaugos).

NGI profile: Lithuania's LitGrid

Getting ready to host the 2011 User Forum



The Lithuanian National Grid Infrastructure (NGI), known as LitGrid, brings together 13 research centres and universities in an ambitious programme sponsored by the Ministry of Science and Education of Lithuania.

With 17 computing clusters and 90 active users, LitGrid is getting ready to co-host the upcoming EGI User Forum, to be held on 11–14 April in Vilnius, Lithuania's capital.

Officially named 'Lithuanian distributed and parallel computing and e-services', LitGrid was established as the country's NGI in 2008, with Vilnius University as the leading partner and now as representative in EGI's Council. Bidding for the event was a natural decision: "It is crucial for us to be involved much deeper in EGI activities and to learn the European experience in grid technology, as well as in similar technologies like cloud computing and HPC," says Algimantas Juozapavicius, LitGrid's project leader.

LitGrid was set up as a project in 2005 "to develop the research computing and communication infrastructure in Lithuania, and to integrate this infrastructure into the emerging European Grid and Baltic Grid infrastructure," Juozapavicius says.

Nowadays, LitGrid operates over 500 processors, with over 30 terabytes storage capacity for the benefit of a growing user base covering numerous research areas. Current scientific applications include modelling of heterogeneous processes in biology, quantum mechanics, material science, analysis and visualisation of multidimensional biomedical data and as-

trophysics. "Usually our users publish about 20–40 scientific articles each year, with data or results obtained with the help of the grid," says Juozapavicius.

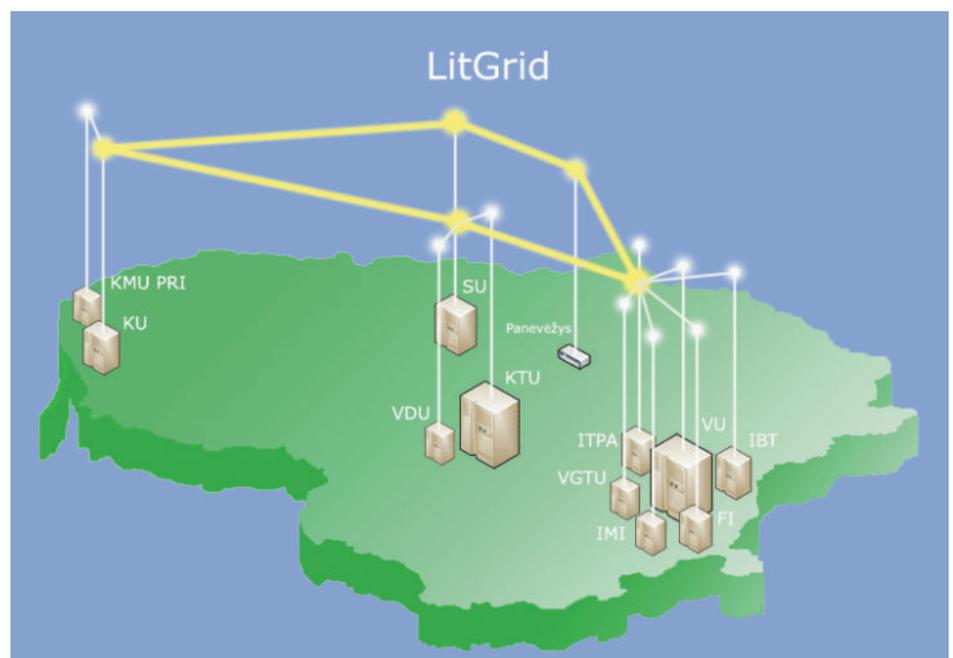
LitGrid is also involved in the analysis, design and test deployment of cloud computing technology in relation to grids and the search for new users never stops.

The Lithuanian infrastructure runs mainly on gLite middleware, but several other operational systems are under consideration, including KnowARC, Globus, UNICORE and CREAM. All LitGrid clusters have EGI certified status and are monitored both globally, with GSTAT, SAM, Stress Test monitors, and locally with Ganglia and Nagios. Lithuania's user support is based on a request tracking workflow managed by a ticketing system.

Despite all the recent progress, LitGrid keeps looking forward: "The aim is to develop the symbiosis of grid-cloud computing, as well as to offer our services to a broader spectrum of society, including academics, public sector and business," he adds.

With this goal in mind, Lithuania is creating a Digital Science and Computing Centre (DiSCC) at Vilnius University. The aim is to enable knowledge transfer and the application of research to business and to foster innovation in digital products.

"DiSCC is expected to be translated into an international centre, attractive for computer scientists, industry, foreign partners and for the state as a customer, helping them to digitise areas of activities by using the most suitable scientific and technological initiatives," says Juozapavicius. •



Distribution of LitGrid's sites in Lithuania

WeNMR: an e-Infrastructure for NMR and structural biology

Sara Coelho profiles the first research community to join EGI

WeNMR is a fast-growing e-Infrastructure community on the frontline of biomedical research, poised to become the first Virtual Research Community to sign a Memorandum of Understanding with the European Grid Infrastructure.

The structural biology community is already a rapidly expanding Virtual Organisation (VO) that accounts for about one in every six life science jobs submitted to the grid infrastructure. In 2010, this totalled around 670,000 jobs, corresponding to about four million CPU hours, on behalf of more than 250 users worldwide.

"Life sciences in general is too broad and the requirements very disparate," says project coordinator Alexandre Bonvin, from the University of Utrecht. The agreement with EGI will allow WeNMR to represent the interests of its community better. "We also hope that recognition by EGI will facilitate the support to our VO from various NGIs [National Grid Infrastructures]," he adds. "This is important for us in terms of sustainability."

The main goal of WeNMR is to offer a user-friendly e-Infrastructure for scientists and to establish a thriving virtual research community in the fields of nuclear magnetic resonance (NMR), small angle X-ray scattering (SAXS) and structural biology, building on the success of the previous eNMR project.

eNMR to WeNMR

While eNMR focused on taking the e-Infrastructure off the ground, "now we are more into operating, consolidating and expanding this e-Infrastructure," says Bonvin. The three-year WeNMR project will provide support to software developers and users and will promote the adoption of e-Infrastructures within the life sciences.

High-resolution NMR allows scientists to measure the distances between atoms of a molecule, which are then used to calculate the three-dimensional structure.

This is not just an academic exercise. Many high-impact diseases, such as Alzheimer's or Parkinson's, are caused by accumulation of protein-type molecules in the brain. Knowing their precise structures is the first step in the search for medical compounds able to defeat the disease.

Calculating molecular structures out of NMR data requires a great deal of computing power and this is where a robust e-Infrastructure can make a difference.

"Many applications are CPU intensive and some are rather complicated to install and maintain," says Bonvin. The CS-Rosetta application, for instance, generates 3D models of proteins but requires 5,000 to 10,000 CPU hours to run. A local cluster of 100 cores would take five days to produce results. "On the grid this can be reduced to less than one day," he adds.

The WeNMR e-Infrastructure uses EGI's grid for the CPU intensive calculations. This will allow small research labs to benefit from state-of-the-art methods even if they do not have adequate local infrastructure resources or expertise.

Bonvin believes that his community, which entered the grid business only rather recently, can bring a "fresh new look at things for the benefit of the wider EGI community."

"We are following a strategy to provide easy and user-friendly access to the grid via web portals," he says. "And we look forward to sharing our experiences and solutions with others in the cadre of the competence centre we are setting up within our VRC". •



"[Joining EGI] is important for us in terms of sustainability."

Alexandre Bonvin,
WeNMR project coordinator

More information:

<http://www.wenmr.eu/>

Project profile: Initiative for Globus in Europe

Promoting Globus in Europe and Europe in Globus

On 20 January, the Initiative for Globus in Europe (IGE) became the first technology provider to sign a Memorandum of Understanding with EGI. Over the next years IGE will contribute new technology components to meet the needs of EGI's users, but what is the Initiative all about?

"Our overarching goal is to help the scientists in the European Research Area," says Helmut Heller, project director of IGE. "We want to do this through coordination of the widespread European Globus development and operation activities."

Globus is a middleware toolkit that allows users to build their own grid, "the same way you use a toolkit with hammer, saw or an axe to build a house," Heller explains. The advantage of this approach is that users can pick and choose which components to deploy. "You don't always have to use the full toolkit: If you know already that you will only do some basic woodworking – building a data grid for example – the saw may be enough."

"Probably the most successful component of Globus is GridFTP," says Heller. This component replaces passwords by certificates, improves high-performance data transfer and is used in other grid middleware, such as gLite and ARC.

The 30-month project was set up in October 2010 to support, develop and promote the Globus toolkit in Europe and to strengthen the influence of European developers within the Globus alliance. It all started with what Heller describes as a true 'grass-roots movement'.

"Before we started we did an informal poll which asked the NGIs and other contact points in European countries if they were using Globus or wanted to do so in the near future,"

he says. The response was overwhelmingly positive and Heller concluded that, perhaps against popular belief, "Globus already has a geographically diverse user base in Europe."

European infrastructure providers such as DEISA and PRACE deploy Globus components in their infrastructures namely Gsissh, which allows interactive access to a remote computers using certificates instead of username or passwords, or GridFTP. "However, many of these Globus users did not know about each other or what the European Globus developers were working on," Heller says. "There was a severe lack of communication and communication channels."

IGE will add a European perspective to Globus: "We will gather and bundle the European requirements and voice them with the core developers in the USA," says Heller. At the same time, IGE will deliver tailored software, training, support and documentation to European research communities.

Another key aim of the project is to increase Globus's visibility in Europe by organising yearly Globus conferences and Globus-themed tracks and workshops in European conferences and EGI's Technical Forums.

The IGE project is barely six months old but Heller is ready to point out the first success. IGE helped a group of users in DEISA to upgrade the Globus tool that submits jobs to the IBM LoadLeveler scheduler from version 4 to version 5, dubbed GT5.

"This enabled the researchers in the Virtual Physiological Human group [VPH] to use supercomputing facilities at the SARA computing centre in the Netherlands where the new GT5 middleware had already been installed," says Heller. VPH can now use computing resources from all over Europe and several sites in the USA – and connecting scientists to computing resources is what we're all here for. •



Steven Newhouse (left) and IGE's Herman Heller shake hands after signing EGI's first MoU.

Gathering requirements from the community

Gergely Sipos introduces the new Request Tracker dashboard

As of January 2011, grid infrastructure users and operations teams can submit their service requirements as tickets to the new Request Tracker (RT) dashboard. The tickets will be handled by the EGI.eu user support and operations teams and categorised by community (for example NGIs, projects, virtual organisations) or according to their status (new, open, accepted, resolved).

The RT system will help the user support team to keep track of what needs to be done and lets the whole EGI community know. It's a global view of what European scientific communities need from EGI.

Requirements will be posted in an open ticketing system, which makes the gathering and solution provisioning processes open and transparent.

This allows the community to check if what they need has already been required by other users, teams or NGI staff and will save time for user communities, for technology providers and avoid duplication of efforts.

The dashboard allows individual users and communities to keep track of what is happening to their ticket and see where their requirement is in the workflow. Users can also comment on existing tickets: you can offer your own solution to the requirements posted on the dashboard as a Technology Provider, an NGI member or member of other EGI communities.

The openness of the new RT system will generate discussions between scientific grid user communities and software developers and catalyse solutions.

The evolution of the European Grid Infrastructure is driven by the users. Therefore capturing and communicating feedback from users to the infrastructure, as well as technology operators and providers, is a key goal for user support and for the EGI collaboration as a whole.

The RT dashboard was born out of this commitment. Previous systems to compile requirements were based for instance on spread sheets. It was time-consuming to manage and it did not promote dialogue between users, support and operations teams.

Instead of reinventing the wheel, the user support team adapted the request ticketing system already in place to prioritise the workload of the EGI-InSPIRE project to collect the community's needs. •

Peter Solagna unveils the operational requirements collection campaign

The Unified Middleware Distribution (UMD) is the integrated set of software components that EGI makes available from Technology Providers within the community. UMD is not static: its functional capabilities will evolve over time in response to feedbacks from users. The requirements can be submitted both by the end users and by operations communities.

Capturing and communicating feedbacks from infrastructure providers to technology providers is a key goal for EGI Operations team, as well as end-users requirements collection is a target of the User Community Support activity. Our main interlocutors in this activity are the National Grid Initiative representatives: they aggregate the requirements collected within their Resource Infrastructure, and submit them through the Request Tickets (RT) system.

Operational requirements are collected over a period of three months, although single requests can be submitted anytime. Middleware requirements are evaluated and prioritised by the Operations Management Board (OMB), before being forwarded to the technology providers.

Requests concerning Operational Tools developed inside EGI are currently collected and they will be analysed and prioritised by the Operations Tools Advisory Group (OTAG).

The current survey focuses on service monitoring and Nagios probes. NGIs can request new probes to be developed, changes to the existing ones, or propose the adaption of new ones locally developed.

Please provide your feedback. Surveys and requirement gathering are crucial if we want the grid middleware development to be a real user-driven process. •

Surveys and requirement gathering are crucial if we want the grid middleware development to be a real user-driven process.

RT dashboard:

<https://rt.egi.eu>

News in Brief

Venus-C Opens Clouds for research, Neasan O'Neill

Venus-C has just launched its Open Call looking for pilot applications to get up and running on the cloud.

Funded by the European Commission, Venus-C brings together industrial partners and researchers to create an enterprise quality cloud service for Europe. To help expand the community they already have, they have launched this new initiative to fund between 10 and 20 new schemes.

These pilots will also help gather additional requirements for the platform, alongside testing and validating it. The projects will have access to all of the resources Venus-C control and be given start-up funds to get the ball rolling.

The experts at Venus-C will also work with the successful applicants to determine what features and capabilities of cloud computing best support their work. They hope to attract interest from a diverse range of disciplines including the Arts & Humanities, Engineering, Health & Life Sciences, Economics, Financial Services, and Natural Sciences. Of particular interest are applications that require dynamic scaling and ubiquitous availability.

Andrea Manieri is one of the coordinators in the engineering group at Venus-C and is looking forward to what this means for the project and

cloud computing in general “Venus-C already has a compelling range of applications but this open call will broaden the scope of the project, and help ensure the future of an academic cloud infrastructure in Europe”

The call is open to public and private research organisations and runs from the 11 January–11 April, 2011. The entire fund is €400,000 which will be equally divided among successful candidates. •

More Information

All call documents are available at: <http://www.venus-c.eu/Pages/OpenCall.aspx>

South-East Europe Regional Operations Centre closes the door, Sara Coelho

After more than five years of work, the South-East Europe Regional Operations Centre (SEE-ROC) has stopped its operations on 31 January and is now decommissioning its tools.

This marks an end to an era for the South East Europe Region that collaborated closely since the first Enabling Grids for E-Science project (EGEE I) in 2004. “It was a fruitful experience that incubated collaboration within the region and spawned off a number of NGIs,” says Kostas Koumantaros, from Greece’s Hellasgrid.

Having multiple National Grid Initiatives instead of one centralised ROC

“allows each NGI to plan individually the upgrades of its infrastructure and possibly minimises administrative effort due to reduced complexity,” explains Koumantaros.

SEE-ROC’s dismantlement follows the closing down of the Central European ROC, which stopped operating on 31 July 2010.

“The most important outcome from the collaboration within the SEE-ROC is the establishment of a vast human network,” says Koumantaros. This potential will be maintained through collaborative Regional Virtual Organisations. •

The Final Countdown:

2010

8 Dec	Cyprus
17 Dec	Georgia
21 Dec	Macedonia

2011

10 Jan	Bosnia Herzegovina
13 Jan	Montenegro
18 Jan	Bulgaria
20 Jan	Israel
31 Jan	SEE-ROC ends

The EGI blog is looking for contributors

The EGI blog was created to be a place to share ideas, thoughts and activities and is part of the on-going developments to improve communication with the rest of the community.

The blog is public and after logging in with your EGI SSO you will be able to comment on any of the material. If you want to contribute to the blog,

please email blog-admin@egi.eu with your SSO username.

The team at EGI.eu have been posting a few entries on the latest developments on the user community support, the LifeWatch project kick-off event and the challenges waiting for us in 2011.

We all welcome your feedback and

we hope your comments will fuel productive online discussions across the many areas of activity within the community. •

The EGI blog

www.egi.eu/blog/

Upcoming events

EGI User Forum: key conference dates

10 January

- > Registration opens

24 January

- > Authors of accepted abstracts are notified

28 February

- > End of early-bird registration fees

1 April

- > End of normal registration fees

Check the user forum website
<http://uf2011.egi.eu/>
for updates to the calendar



Booth registration
open until **1 April**

See you in
Vilnius !

Enabling Communities at the TERENA Networking Conference, Carrie Solomon

Specialists from major European research networking organisations, research projects, universities and industry will come together at this year's TERENA Networking Conference (TNC2011) from 16–19 May, 2011 in Prague, Czech Republic.

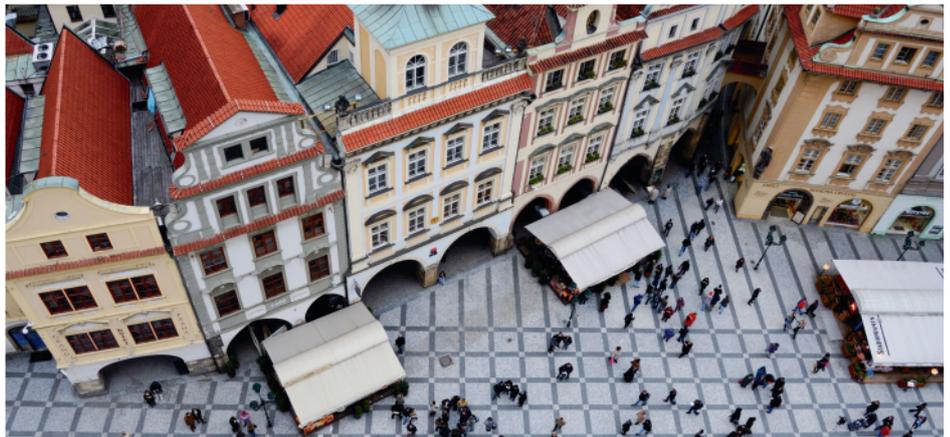
Keynote speeches and interactive discussions between experts will bring to life the theme 'enabling communities', and provide participants with an overview of the latest developments in research networking, both in the technical field and in the areas of application and management.

Morning plenary sessions will set the tone for topics explored in more detail throughout four parallel tracks: 'future networks', 'supporting collaborations and advanced applications', 'communities and identities', and 'digital lifestyle'.

Other conference highlights include:

- > Ian Bird, WLCG (World LHC Computing Grid) project leader, will talk about the Large Hadron Collider (LHC), the status of the project, the advanced applications and infrastructure that support it, and how research and education networks can help make large-scale, global collaborations in science successful.

- > Jaroslav Koča, Scientific Director for Life Sciences, Central European



Institute of Technology (CEITEC) will describe how networks enable his work in biomedical research, one of the most important areas of global scientific collaboration.

- > Connections between high-demand research, the 'digital native' student and the everyday work of the research networking community will be addressed throughout the conference in discussions about the future Internet, new ways to expand the use of identity federations, grids & clouds, virtual organisations & collaborations, e-science ecosystems and much more.

Registration for the conference opens on 21 February and the fee includes access to all conference sessions, materials and social events.

You can participate in a lightning talk – a five-minute presentation focusing, for example, on an idea, a successful project, or an invitation to collaborate – or submit a poster presentation. A limited number of free admissions for students are also available for winning poster submissions. •

More Information

TNC 2011 webpage:
<http://tnc2011.terena.org/>

Gyongyi Horvath:
horvath@terena.org