

Inspired

ISSUE 13
OCTOBER 2013

news from the EGI community

(Juergen goevert / wikicommons)



TOP STORIES

The EGI Strategy for 2020

page 2

Q&A: FedCloud commercial provider

page 3

ENM - a BioVeL FedCloud use case

page 5

EGI-EUDAT-EISCAT_3D collaboration

page 7

Promoting Desktop Grids in EGI

page 8

MORE

- 02 Achievements of the EGI - HelixNebula collaboration
- 09 Steven Newhouse says farewell to the European grid
- 10 The CHAIN-REDS Science Gateway
- 11 New features of AppDB
- 12 Kick-starting the EGI Resource Allocation Pool
- 12 A new Service Management certification



**European Grid
Infrastructure**

www.egi.eu

This Issue

- > Luděk Matyska announces the EGI Strategy for 2020
- > Małgorzata Krakowian and Gergely Sipos write about the EISCAT_3D collaboration
- > Robert Lovas makes the case for Desktop Grids
- > Nuno Ferreira introduces a FedCloud use case
- > Michel Drescher leads a Q&A with David Blundell from 100 percent IT
- > Steven Newhouse says farewell after five years of European Grid
- > William Karageorgos and Marios Chatziangelou give their AppDB update
- > Roberto Barbera and Diego Scardaci introduce the CHAIN-REDS Science Gateway

Plus:

- > Kick-starting the EGI Resource Allocation Pool
- > Achievements of the EGI-Helix Nebula collaboration
- > A new certification in Service Management
- > And our upcoming events

Comments and feedback to: sara.coelho@egi.eu



This issue's cover image is a cardinal fish (*Pterapogon kauderni*) and a reference to a certain Douglas Adams novel. (Illustration: Juergen goevert / wikicommons)

Achievements of the EGI-Helix Nebula collaboration

Carmela Asero

Recent developments in the EGI Federated Cloud highlight the high potential of EGI contribution to the Helix Nebula initiative for a scientific cloud in Europe.

The EGI Federated Cloud's latest achievements were under the spotlight in September at the EGI Technical Forum in Madrid and at the 3rd Helix Nebula General Assembly in Heidelberg. The proof-of-concept from the European Space Agency (ESA), previously run on Helix Nebula's commercial providers, was repeated using EGI's Federated Cloud academic resources during the Technical Forum. The results of the successful experiment were demonstrated live by Salvatore Pinto, EGI.eu's cloud technologist, who showed promising interoperability results with the use of an OCCI connector for SlipStream.

The SlipStream software was developed by SixSq and has been tested by the Helix Nebula project as a technology solution to broker between the user and the different clouds on offer, both commercial and academic. The demonstration showed how academic and commercial cloud resources can be successfully run alongside each other and was also an example of how the EGI community promotes innovation by partnering up with commercial technology providers such as SixSq. These results prompted very positive feedback at the 3rd Helix Nebula General Assembly in Heidelberg and work is already in progress for the next interoperability milestone. This will be the CERN proof-of-concept on a hybrid cloud environment built from a pool

of publicly funded resources provided by the EGI Federated Cloud and from Helix Nebula's commercial partners. The CERN proof-of-concept is expected to be deployed in October 2013. EGI.eu and Helix Nebula are also reinforcing their cooperation through an agreement currently under negotiation. The MoU will cover joint activities in key areas such as interoperability, deployment and communications.

More information

EGI FedCloud
<http://go.egi.eu/cloud>

HelixNebula
<http://helix-nebula.eu/>

SixSq <http://sixsq.com>

The EGI Strategy for Horizon 2020

Luděk Matyska, chair of the EGI Council, writes about EGI's future as a secure, federated data-analysis infrastructure for the European Research Area

The European Research Area (ERA) needs to support multidisciplinary research collaborations that cross national and intellectual borders to tackle society's grand challenges. To produce the excellent science that will deliver innovation and growth, Europe's researchers need easy-to-use integrated services providing access to high-capacity and high-quality computing and storage resources, wherever the resources and the researcher are located.

This is what EGI aims to become. In June 2011, EGI adopted a common strategy document 'Seeking New Horizon – EGI's Role in 2020' to set the path for future evolution of the services to serve the changing needs of the ERA. The paper was updated by the EGI Council in September 2013, to reflect the latest evolution of the community and the inputs from a number of policy papers issued in the last two years.

The EGI Collaboration (through the EGI Council), composed of representatives from the National Grid Initiatives (NGIs) and supported research communities, is now proposing an ambitious investment in open distributed computing and data infrastructures that will build on the existing European and national investments to:

> **Provide enabling services to researchers** by adopting an expandable and user-centric service portfolio to retain its current research communities and attract new ones.

> **Operate an unprecedented European capability for high-throughput data analysis** that expands EGI's current Federated Cloud infrastructure to 10 million computing cores and 1 Exabyte of storage by 2020. EGI will build on its current collaborative resource allocation model based on the resources coming from the NGIs and other organisations to also support peer-reviewed or pay-for-use access.

> **Provide flexible Virtual Research Environments** that simplify access to EGI's resources and make it easier for researchers to work and make the most of the NGI expertise to introduce technical innovations into production across Europe.

> **Identify and develop EGI's Human Capital** to upskill research communities by establishing national centres of excellence within the NGIs to promote knowledge transfer from the EGI community to tomorrow's data scientists.

Collectively, these activities will stimulate technology and result in research innovations through effective use of high-throughput data analytics on a high-capacity and integrated European infrastructure by research communities in the ERA.

Impact for the researchers and research communities within the ERA

> **Excellent science** through innovative services and their continuous improvement from both public sector or commercial providers;

> **Every researcher digital** through the EGI distributed competence centres;

> **A scalable computational and storage research infrastructure** through the integration and federation of resources from both public and private sectors.

Impact for the NGIs and other organisations participating in EGI.eu

> **A growing and developing a national e-Infrastructure** through EGI.eu as a focal point for integration and collaboration at a European level;

> **NGI Virtual Centres of Excellence** through maximising the collective knowledge and skills developed within the EGI collaboration.

Impact for the European Commission and European citizens

> **A European coordination and governance** through the organisations representing member states and research communities as stakeholders in EGI.eu

> **A distributed open computing and data infrastructure for Europe** through the operational integration of distributed heterogeneous resources

> **High utilisation of public sector investments** through pooling national and EC investments into an integrated initiative accessible to the ERA.

More information

The EGI Strategy for 2020

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

Ecological Niche Modelling running on the EGI Federated Cloud

Nuno Ferreira writes about a pioneering use case of the EGI Federated Cloud and its scientific applications

The EGI Federated Cloud is a new type of e-infrastructure to be rolled into production next Spring. In this development phase, use cases from research communities are key to steer and fine tune its capabilities. One of these use cases comes from **BioVeL** - an FP7 project set up to provide a virtual e-laboratory for biodiversity researchers.

Ecological Niche Modelling

BioVeL's goal is to give users a customised environment with readily available solutions to run complex workflows, using data and web service resources from a wide range of sources. One of the most robust and popular workflows used by the biodiversity community is Ecological Niche Modelling (ENM). The **ENM workflow** uses algorithms to model the theoretical distribution of




Earthworms (pictured: species *Octolasion cyaneum*) play a crucial role in terrestrial ecosystems. They help to break down organic matter into compost and they improve fertility of arable lands. This is why scientists use Environmental Niche Modelling software, such as openModeller, to find out which environmental variables are likely to impact earthworms, and how will their populations evolve over time.

species according to known occurrences in the field. The maps generated through this process represent the potential distribution of a species in a particular geographic region at a particular time. Predicting species' distributions and understanding why they change is important in many research areas such as biology, ecology and evolution.

ENM models, such as the earthworm distribution presented in the box, are created by **openModeller (oM)** – a flexible cross-platform environment, able to run multiple algorithms. OpenModeller was originally developed by the CRIA institute (Reference Center on Environmental Information) in 2003 and, over the years has become increasingly popular. This high-demand has led BioVeL to join forces with EGI to setup and operate an oM web service instance in Europe.

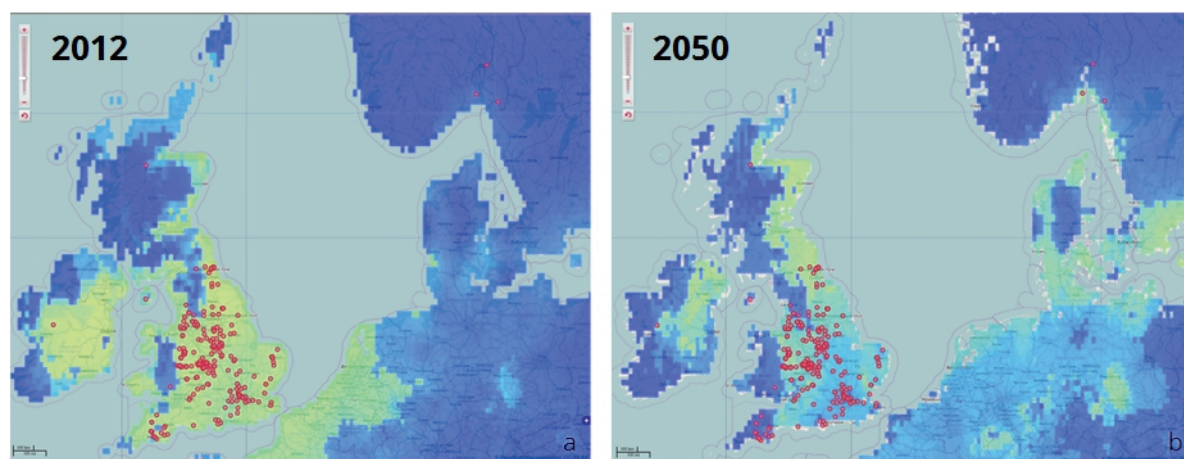
The EGI Federated Cloud

The European Grid Infrastructure was originally set up to federate computer resources via grid middleware. But as scientific computing paradigms evolve, so EGI needs to adapt to the new trends. The EGI Federated Cloud was born from a concerted effort between resource providers (RPs), technology providers and user communities willing to pioneer the cloud federated model. The federation layer for this new research e-infrastructure is based on open standards for sustainability purposes. Furthermore, the federation is agnostic to the cloud technology deployed at each RP, as long as standard interfaces are exposed to the federation in a seamless way. The federated cloud is moving towards production level at a fast pace, as the services become fully operational and integrated into the EGI



(cont.)

Environmental Niche Modelling in action - distribution of earthworms over time



Octolasion cyaneum is an earthworm species important for soil quality. The projections obtained from the Environmental Niche Modelling workflow show that the suitable habitat of the species is likely to shift towards North/Northeast in the next 40 years. The colour codes indicate habitat suitability, from blue (not suitable) to green/yellow (suitable) as calculated by the model for today (a) and for 2050 (b). The red dots indicate recorded occurrences of the species.

Many thanks to Matthias Obst, leader of the BioVeL's Work Package 3 (Services for biodiversity analysis) for providing this example.

production infrastructure and resource centres are certified.

A trinity

After focusing on the requirements of the biodiversity scientific community to make the openModeller web service available for ecological niche modelling, a solution was found with the following actors:

- > **BioVeL**, providing the use case
- > **EUBrazilOpenBio**, providing the ENM as a service
- > **EGI Federated Cloud**, providing the computing resources.

From the user perspective, there is no need to change the routine scientific methodology to achieve the expected results, while running the ENM workflow. Everything happens under the hood of a complex set of technologies. The user simply has to send a request from his client interface to receive the results.

The openModeller extended web service developed by EUBrazilOpenBio, orchestrates

the user requests and takes care of dispatching the calls to the back-ends in the cloud federation, via VENUS-C COMPSs middleware. This set of technologies allow for capabilities such as brokering and load-balancing with dynamic deployment of virtual instances - a must have for services requiring high-availability and multi-user calls.

Resource providers from Spain and Czech Republic were targeted for the implementation of the use case. The cloudified openModeller application proved to be reliable when in conjunction with the COMPS framework.

The benefits for BioVeL

As a result of this concerted effort, the EGI Federated Cloud is now able to offer BioVeL:

- > An instance of the popular openModeller application is now available in Europe to serve the biodiversity community
- > Dedicated support from cloud experts to optimise the

application virtual machine

- > Technology able to scale available resources based on the load of user calls
- > Cloud implementation model to cloudify other biodiversity relevant web services
- > Sustainability of a popular web service.

More information

The EUBrazilOpenBio initiative has developed an IT infrastructure offering resources (e.g. computing capacity, software) as a service, aggregating taxonomic, biodiversity and climate data, with an easy-to-use Virtual Research Environment.

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

BioVeL project
<http://www.biovel.eu/>

EGI Federated Cloud
<http://go.egi.eu/cloud>

Q&A: 100 Percent IT, a FedCloud commercial partner

Michel Drescher speaks to David Blundell, Managing Director of 100 Percent IT - the first commercial provider to partner with the EGI Federated Cloud

1) Tell us a bit about 100 Percent IT

100 Percent IT is a business Internet Service Provider based in the UK. We've been established for over 13 years and we specialise in providing high availability cloud / virtualisation services and networks. We enjoy working with partners where the benefits we offer - speed, resilience, data security and ease of use - are important.

2) How did you get to learn about the EGI Federated Cloud activity?

I first heard about the EGI when I met David Wallom of Oxford University at an STFC event at Harwell in April this year. A lot has happened since then!

3) What were your reasons for starting to collaborate, and then joining?

Initially we were interested in the intellectual challenge of integrating our system in to the

EGI. We knew we'd developed a robust automated system based on OpenStack but this was a great opportunity to prove our technology and systems with world class users expecting production-strength systems.

We are very pleased that the feedback so far has been extremely positive, and we are looking forward to working with many more of you!

4) How do you compare your certification experience with what you expected? Which are your recommendations for improvements?

A technically challenging certification process gives EGI users confidence that all the providers on the EGI framework have a high-quality product. The EGI documentation had a few gaps which delayed the integration work, but we worked with very helpful engineers across Europe to solve the issues. So, we enjoyed the challenges of

integrating; particularly developing our own automated installer that means we can roll out an EGI cloud on bare metal in just a few minutes. Penetration testing to ensure the security of all providers systems would be an added benefit that the EGI could offer to resource users.

5) Open markets frequently provide consumers with many different offers that are difficult to compare. Which instruments would you recommend to ensure that users can make a better choice?

We believe in designing resilient, reliable systems as standard and we offer a 100% SLA with 100x service credits. But SLAs are only part of the picture. Sometimes the compensation you may receive after an outage is not as relevant as access to your data when you need it. What users need to know is how much uptime they can expect and how quickly problems are resolved. EGI makes this transparent by publishing real time statistics for each provider so that users can choose from a suitable provider in the league table.

6) The EGI Federated Cloud will go into production soon. How will this affect your business planning?

We are actively working to move the 100 Percent IT cloud into the Federated Cloud production site as you read



100%IT

"We knew we'd developed a robust automated system (...) but this was a great opportunity to prove our technology and systems with world class users expecting production-strength systems."

David Blundell,
Managing Director
100 Percent IT



this. Over the past couple of years we've invested a lot of time and money developing what we believe is a technically excellent solution and are looking forward to it being in production within EGI shortly.

100 Percent IT is planning rapid growth as adoption of our cloud solution progresses and we hope that EGI will be an important part of our future.

7) In your opinion, which federation policy would apply best? Do you see benefits in a central resource provisioning model; or rather stay with pure bilateral agreements?

100 Percent IT always wants to be a supportive partner to the EGI and work with you to develop systems that work well for all parties.

A central resource provisioning system could make access easier for end-users so 100 Percent IT would like to facilitate this and as an example would be willing to commit e.g. 20% surplus resources over and above those in use at any given time to the central pool.

8) A federation is always a give-and-take situation for commercial providers. A clear benefit is an increase in market penetration. Can you name five federation services you would abide to become a member of the federation?

1. A single API / authentication mechanism for all end-users. This means that providers only need to configure one way of connecting into their system to gain access to multiple larger contracts.
2. Centralised monitoring. This enables providers to be compared on a like for like basis and will enable technically competent companies regardless of size to shine.
3. First line support. Support queries being answered by local help desks enables providers to focus on development while being available as back up should a complex issue arise.
4. Computer Security Incident Response Team. A security

vulnerability discovered at one provider can be quickly analysed and the fix applied across all providers increasing security for all.

5. Future development. Providers can work together with the EGI to develop improvements that benefit the whole community, ultimately increasing usage across the board and creating more resilient systems. An example of this is the centralised VM image distribution system where VM templates are created centrally and added to each provider's system allowing end users to choose the template that suits them from any provider.

More information

100 Percent IT
<http://www.100percentit.com/>

EGI Federated Cloud
<http://go.egi.eu/cloud>

EGI – EUDAT – EISCAT_3D collaboration

Małgorzata Krakowian and Gergely Sipos on the work that will help EISCAT_3D to address a big data challenge

The EISCAT association operates the world's largest system of incoherent scatter radar installations, used to observe the high-latitude atmosphere and ionosphere.

The next generation of radar systems – **EISCAT_3D** – opens new scientific opportunities, but it also introduces significant computational challenges. During its first operation stage starting in 2018, EISCAT_3D will produce 5 PB data per year, and the total data volume will rise up to 40 PB per year when it hits full operational status in 2023. How will these large volumes of experimental data be handled?

These needs are an excellent example of a 'big data problem', whose size is beyond the capabilities of current database and software technology.

This is where EGI can bring added value.

Towards a Big Data strategy for EISCAT_3D

Together with EUDAT, EGI established a 'study case' partnership with EISCAT_3D under the umbrella of the ENVRI project to identify solutions for EISCAT_3D's computational challenges. These solutions are expected to be reusable by other ESFRIs in the ENVRI cluster of environmental sciences projects.

The collaboration aims to:

- > Analyse the EISCAT_3D data infrastructure to capture the requirements of the scientific community for applications to work and process data.
- > Analyse EGI and EUDAT service portfolios to identify what could be offered to fulfill EISCAT_3D's big data requirements.

An artist's impression of the EISCAT_3D array (www.eiscat3d.se)



- > Advise EISCAT_3D on the setup of a strategy and an infrastructure for big data.

What EISCAT-3D needs

The collaboration started with capturing data management requirements through two online surveys targeted at the EISCAT_3D community. The activity began in August, and the feedback collected so far reveals a need for a set of services to support every stage of the data processing pipelines of the EISCAT_3D community:

- > EISCAT_3D will need new types of visualisation products, for example to represent spatial distributions of plasma parameters in 3D.
- > In the current EISCAT archive, data can only be found by observation date and radar name. To support complex queries, they need to setup a metadata schema and specific data descriptions. Interfaces and engines for data description, curation and lookup are also of prime importance.
- > A method for 'flicking' through short sequences of plots to see how the data evolves through time would be useful. This data will be organised by time, and possibly cached at a user-level to enable interactive plotting applications.
- > EISCAT_3D needs data export services to sharing of workflows

among collaborators. The transport method (e.g. ftp, http, smtp) should be adapted to the existing tools of the community.

- > A need for data sharing facilities that allow the mounting of data from online read-only resources through standard file systems (sshfs, nfs or similar) into local applications.

In summary, there is a need not only for new technologies, but also for new ways to work collaboratively and share data as part of their research workflow. What EGI and EUDAT can do is provide the tools and technologies to integrate existing and future applications from the EISCAT community with European resources.

The surveys remain open until the end of the year to capture additional use cases. This will enable the collaboration to define the services that need to be assembled and developed for EISCAT_3D to address its 'big data challenge'.

More information

If you wish to get involved in the definition of the EISCAT_3D big data strategy, get in touch directly with Małgorzata - operations@egi.eu

More computing resources at no cost: Desktop Grids

Robert Lovas invites NGIs to join a Virtual Team to promote Desktop Grids in EGI

What are Desktop Grids?

Desktop Grids pool the unused capacity of personal computers for scientific purposes. Because the resources are donated by private citizens, Desktop Grids are able to provide a cost-efficient and large-scale capacity, which is also self-financing, self-updating and self-maintaining.

With this technology, NGIs will be able to access a resource pool offered by volunteers and organisations worldwide.

What EGI members can do with Desktop Grids?

Desktop Grids are already available for NGIs and VOs thanks to the long-term collaboration between the EGI community, MTA SZTAKI from the Hungarian NGI, and other key members of the International Desktop Grid Federation. The result was a bridging technology that provides interoperability between the EGI core infrastructure and Desktop Grid solution at job management, monitoring, and accounting levels.

Why should you consider Desktop Grid resources?

> **Integration** - seamless integration between EGI services and Desktop Grids means that same jobs, inputs, parameters, and submission tools can be used by the scientists thanks to virtualisation technologies

> **Visibility** - the scientific and technical goals and achievements of your NGI/VO will reach thousands of citizens, with further opportunities for stronger engagement with the volunteers and the society in general

> **Sustainability** - involving a worldwide community of individual resource providers will

contribute to your sustainability or Social Responsibility Plan

What will you get?

The most important benefit is extra computing resources. MTA SZTAKI operates two large-scale volunteer Desktop Grids for EGI users through production-level bridges:

> **SZTAKI Desktop Grid** dedicated to NGI_HU users, comprises 100,000 registered computers with 600 million normalized CPU hours delivered since 2005.

> **EDGeS@home** (soon to be renamed EGI@home) is allocated for EGI users and has 24,000 registered computers and has delivered about 300 million normalised CPU hours since 2009.

Interested? A Virtual Team is the next step

We would like to encourage NGIs and VOs to join a new Virtual Team to promote Desktop Grid Solutions (DG VT), where we will:

1. Promote Desktop Grid-related technologies in the EGI community
2. Promote the use of available bridged Desktop Grid resources (e.g. EDGeS@home) by more VOs
3. Increase the number of heavily-used EGI applications on the integrated Desktop Grid infrastructure with focus on widespread tools, solutions and approaches
4. Improve existing documentation with updated roadmaps, training materials, and manuals.

5. Complete the final remaining steps for full integration, concerning for example support tools
6. Find joint EGI/IDGF champions

The Virtual Team

The VT will support NGIs and VOs that wish to try running their current applications in this



An idle machine wonders what to do with its spare time...

combined infrastructure, and if necessary fine-tune them taking into consideration the nature of Desktop Grids.

The immediate goal of the VT is to increase the size of the community who speaks not only Grid but also Desktop Grid, and support it with more interconnected VOs, more applications and better materials.

In the long-term, we hope to increase the number of new user communities who start using EGI thanks to the alternative approach of Desktop Grids.

More information

> The VT will be supported by the International Desktop Grid Federation Support Project <http://idgf-sp.eu>

> To join the VT and learn more about Desktop Grids or: <http://go.egi.eu/vtdg>, or email: robert.lovas@sztaki.mta.hu

> Robert Lovas is the NGI International Liaison for Hungary

So long and thanks for all the.... air miles

Steven Newhouse says farewell after five years of European Grid

Five years ago (almost to the day) I returned to CERN after a break of many years to join the EGEE-III project as Technical Director. Back then, the Higgs had not been discovered, the LHC was shutdown awaiting repair and the European Grid Infrastructure was still a pipe dream.

The next five years were a roller coaster journey across Geneva, Amsterdam, Brussels and many other cities, airports and hotels in Europe and around the world. Together, we completed the last of the EGEE projects and transitioned to the EGI-InSPIRE project and its novel operating and governance model based on national structures, brought together by a central coordinating body. This represented a drastic change in our community's structure.

The result was a success and this achievement has quite rightly gained considerable attention within the e-Infrastructure community: many of the ESFRI projects, for example, are now considering the option of a central coordinating body for their research infrastructures.

Where are we now?

EGI.eu, the coordinating body established as a result of the European Grid Initiative Design Study project, has been operating for over three years. The Higgs boson has been discovered at the LHC, and the Nobel Prize for Physics was awarded just a few weeks ago. The current capacity and capability of the infrastructure dwarves peer initiatives around the world, and includes HPC,

grid and cloud resources, running with many different technologies, such as gLite, ARC, Unicore, Globus and QCG.

We are able to federate and operate resources, to provide a high-throughput or cloud-based analysis infrastructure, to make the most of our human networks and to use their expertise to develop new innovations across new and existing research communities.

These are achievements we can all be very proud of – and could only have been achieved with the support and dedication of the staff within EGI-InSPIRE and across the supporting national and European initiatives.

Where does EGI go next?

Within the broader public sector and specifically within the European Research Area, EGI.eu and EGI have the opportunity to help define the e-infrastructures of the future. Collaboration and sharing will remain an important aspect of this work, but increasingly the need to focus on consuming services from a range of different providers – in both the commercial and public sectors – will grow. This in my view means an evolution of EGI towards a sustainable

infrastructure.

A sustainable infrastructure can only be built on sustainable foundations. Working directly with service providers within the larger resource centres across Europe, regardless as to whether they sit within an NGI, an EIRO lab, or in the commercial sector, would allow EGI.eu to focus on the quality, compatibility and integration aspects of federating services for diverse research communities. Around this service offering, EGI.eu and the NGIs can build up the knowledge and expertise needed to help the next generation of data scientists across Europe to navigate through the data deluge and turn it into the innovation that will meet society's needs.

And finally...

It has been a wonderful experience over the last five years meeting people all around the world and seeing the enthusiasm and commitment that they have for distributed computing and data infrastructures, and EGI in particular. I firmly believe that EGI and EGI.eu are capable of being part of this future if it is ready to continue to adapt and evolve to new challenges the way it has in the past.



CHAIN-REDS Science Gateway: linking up HPC, cloud and grid resources

Roberto Barbera and Diego Scardaci introduce a new way to make the most of DCIs

The CHAIN-REDS project has deployed the CHAIN-REDS Science Gateway as a portal to interface different types of Distributed Computing Infrastructures (DCIs): local cluster, HPC, cloud and grid.

The portal, built with the Catania Science Gateway Framework, was presented at the EGI Technical Forum in Madrid, in a demo entitled *Managing and using interoperable DCIs through a standard-based Science Gateway*.

The demonstration of the CHAIN-REDS Science Gateway showed the interoperability and inter-operation among different DCIs, including clouds using OCCI and SAGA as standard interfaces, and the CHAIN-REDS Science Gateway as Virtual Research Environment. The vision behind the development of the portal is that standard-based interoperability should be enabled not only across middleware but also, and more importantly, across computing paradigms (e.g. grid, cloud). As a true Virtual Research Environment, the CHAIN-REDS Science Gateway was designed with an intuitive graphic user interface that includes point & click and drag & drop functionalities.

In practice, the portal allows:

- > A scientist to run applications seamlessly on HPC machines, grids and clouds. This is important because each DCI has its strengths and researchers may wish to combine different added values for a particular project.
- > The cloud tenant of a real or virtual organisation to easily

manage cloud resources pledged by providers based on different middleware stacks.

The demonstration in Madrid mapped these two aspects of the vision on two corresponding use cases:

- > In the first case, the demo showed how a researcher, or another end-user, can sign in on the CHAIN-REDS Science Gateway using his/her federated credentials, select an application from a menu and transparently execute it on HPC machines, grids and clouds.

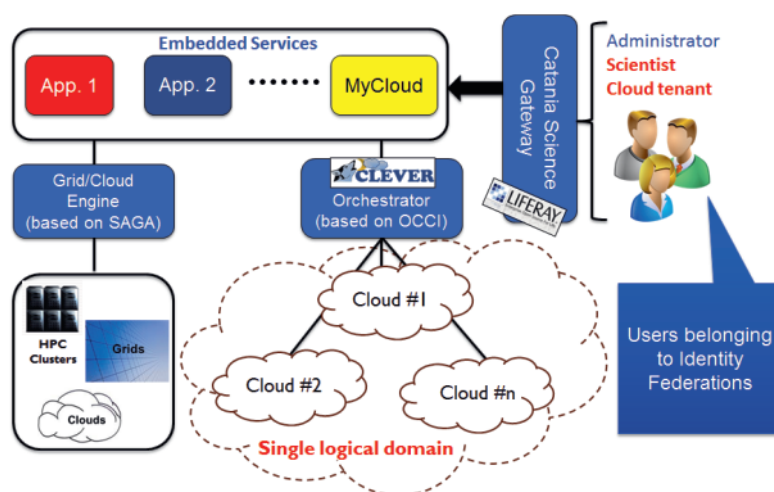
- > In the second case, a cloud tenant signed in (again using federated credentials), selected virtual machines from a geographically shared repository and deployed/moved/copied them across the 'personal' multi-cloud that they manage. Using dynamic DNS functionalities, the virtual machines were assigned to the same domain name, independently of the site where they were instantiated and regardless of the underlying cloud middleware stack.

Both perspectives were demonstrated on a multi-middleware

production infrastructure that included sites belonging to the European Grid Infrastructure, the EGI Federated Cloud and to other cloud infrastructures. In total, eight cloud sites (one from an Egyptian SME) from six countries, running three different stacks, were integrated in the testbed.

Now that the feasibility of the CHAIN-REDS vision has been successfully demonstrated, the next steps are to:

- > Foster the deployment of more cloud infrastructures in the regions addressed by CHAIN-REDS to widen the testbed both in size and geographic coverage;
- > Promote the EGI Federated Cloud model and the extension of its infrastructure to other regions of the world in order to support global VRCs.



More information

CHAIN-REDS project
<http://www.chain-project.eu>

CHAIN-REDS Science Gateway
<http://science-gateway.chain-project.eu/>

New features in AppDB – better search, new VM Marketplace

William Karageorgos and Marios Chatziangelou give an update on a major overhaul

The EGI Applications Database (AppDB) is a central service that stores scientific applications, ready to be used by researchers of all fields of science. The AppDB is always evolving and this article summarises the newest features.

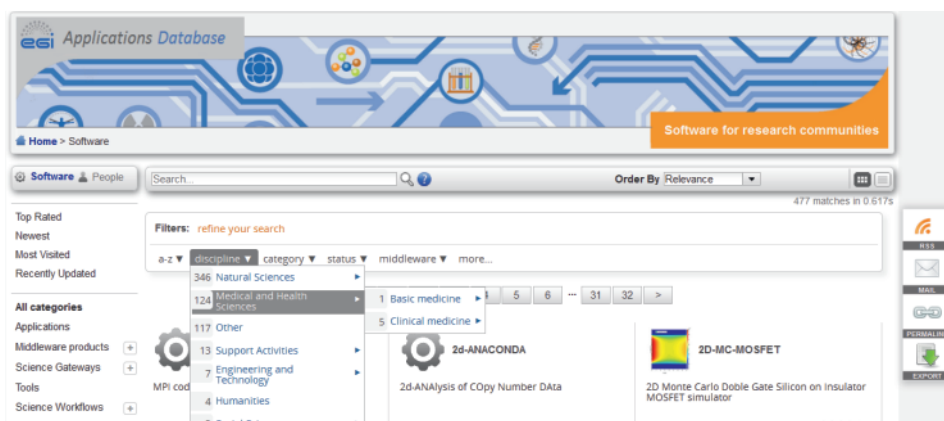
Search performance

The first of these changes is all about optimising AppDB's content for the biggest search engine, Google. The problem, up to recently, was that search engines had, in general, a hard time indexing AppDB's content about software.

The solution to this was to implement **Search Engine Optimisation (SEO)**. AppDB version 4.3.0 adjusts the meta tag section of dynamically-produced HTML content in order to offer details about software titles, authors, owners, ratings, descriptions, abstracts, or keywords. Meta tags have also been extended with Open Graph properties and content now includes web semantics following the microdata specification, allowing Google to create detailed snippets on its search results pages. All dynamic metadata gets posted daily to Google to keep search results up to date.

Result? An immediate increase of AppDB's ranking in search results.

We have also introduced a **faceted search mechanism**, which binds syntactical and tag search systems in a versatile graphical way. Faceted searches



A screenshot of AppDB, showing the faceted search tool (dropdown menus) and the social toolbar on the right.

allow users to retrieve search before performing the query. Each time a user modifies the selection of attributes, the results are refreshed and a number next to each attribute is displayed in order to indicate the count of remaining entities that match.

New science categories

Up to recently, AppDB had been using the old EGEE classification of software into a few simple top-level domains that reflected usage rather than scientific diversity. As part of a community-effort, the classification system was rewritten to allow for a comprehensive list of top-level disciplines with little or no overlapping, and a well-defined set of subdisciplines, up to three levels deep. The AppDB existing software entries have been re-mapped to the new system, providing better statistics on usage across disciplines.

Social tools

AppDB now has a social toolbar to help users advertise their work to their respective audiences. The toolbar appears

in each software details view, and allows readers to post entries about AppDB in any of the major social networking sites.

Virtual Machine Marketplace

Keeping up with the times, AppDB is about to introduce another major feature - a VM Marketplace to support Virtual Appliances.

Virtual Appliances are virtual machines packaged with an operating system and software applications, ready to be used with minimal or no set-up needed. The VM Marketplace is ready to roll into production when the EGI Federated Cloud goes live in 2014. The marketplace will be the place to create, publish, and archive Virtual Appliances.

Feedback is important!

AppDB - <http://appdb.egi.eu>

Your feedback is much appreciated - we will be glad to hear from you on <http://go.egi.eu/appfb>.

A new certification in Service Management



Sy Holsinger

Sixteen members of the EGI community received certification in "Foundation in Service Management for federated IT infrastructures" according to the new FitSM standard.

The first FitSM certification course in lightweight service management for federated IT infrastructures was held at the Technical Forum in Madrid. The newly established standard - FitSM - is backed by TÜV SÜD, an international certification authority. FitSM introduces the fundamentals of service management processes and methods, focusing on the challenges of managing IT services across federated communities.

A total of sixteen participants, including five EGI.eu staff members, took part on a one day intensive course and exam. The training and certification programme was designed to develop

community knowledge and skills in IT Service Management.

"I found the training course very enjoyable. The trainers knowledgeably introduced the aims and methods of IT service management as well as the benefits and limitations both in general and in the context of federated e-Infrastructures. I left the course with many new ideas for organising my work," said Christian Sættrup, Grid Operations at the Niels Bohr Institute.

The FitSM standard, developed by the FedSM project in collaboration with client partners EGI.eu and the Polish and Finnish National Grid Infrastructures, is lightweight and tailored for federated environments, such as EGI.

FitSM provides a baseline level that can support 'management interoperability' in federated environments where disparate or competing organisations

must cooperate to manage services.

"We were really delighted with how much interest was shown in the training - it shows the desire for ITSM knowledge in the EGI community," said Thomas Schaaf, FedSM Project Coordinator.

"It was also great to see that all the participants passed the exam and achieved the FitSM Foundation Certificate from TÜV SÜD."

More information

To learn more about FitSM and training opportunities, contact training@fedsm.eu or visit <http://www.fedsm.eu/fitsm>

To learn more about EGI.eu's experience with implementing service management, contact policy@egi.eu

Kick-starting the EGI Resource Allocation Pool with 1 million CPU hours and 70 Tb of storage

Peter Solagna

In June 2013, the EGI Council approved the implementation of a EGI Resource Allocation Pool to support resource needs of user communities.

The Resource Pool will be available for new user communities approaching the EGI infrastructure for the first time, or for new, small-scale proof-of-concept experiments run by the communities already working with EGI, and will be offered with a public call for interest.

The first step to implement the idea was a call for contributors

to the centralised Resource Pool. To that effect, EGI launched a survey among the NGIs, who can collect offers at national level, and the Resource Centres, asking them to contribute a subset of their resources to the pool.

Sites could contribute by offering:

- > Guaranteed resources, reserved exclusively to a specific VO and allocated for a specific use case; the computing activity can start immediately.

- > Opportunistic resources, open to ad hoc users but subject to local availability; users may have

to wait for unused capacity to run their jobs.

As a result of the call, six resource centres are currently offering a total of 1 million guaranteed CPU hours, distributed across about 700 cores, to the Resource Pool.

On top of this, users will be able to access additional CPU time from the sites that offered opportunistic usage of their capacity.

The storage capacity offered in the resource pool is about 70 TB, shared between guaranteed space and opportunistic use.

Upcoming events

EGI towards Horizon 2020 Workshop

Tiziana Ferrari, the new EGI-InSPIRE project director, is organising a workshop to engage actively with the EGI communities and develop ideas towards the Horizon 2020 work programme. The meeting will be held in Amsterdam, between 4-6 December.

The workshop aims to:

- > present the EGI strategy to the whole community;
- > discuss the Horizon 2020 Work Programme 2014-2015;
- > collect and discuss plans and strategies in a collaborative environment.

The workshop will focus on ideas to allow EGI to develop

new services and business models, the integration and support of multidisciplinary science, the human networks and the set up of a EGI distributed competence centre.

The EGI Federated Cloud will also be a major theme during the meeting that will explore how to extend its capabilities and its supported user communities.

The workshop is intended for but is not limited to:

- > Representatives of existing and new user communities
- > Research Infrastructure representatives
- > NGIs and EIROs
- > Technology Providers

Expected outcome

The workshop will capture stakeholder's ideas for new EGI InSPIRE mini-projects and, most significantly in the longer term, for innovative project collaborations that match the Horizon 2020 vision.

More information

Registration and programme information are available in Indico:
<http://go.egi.eu/H2020>

EGI Community Forum 2014

The EGI Community Forum 2014 will take place on 19-23 May, in Helsinki, Finland.

The event will bring together over 400 researchers, engineers, decision makers and IT people to discuss the use and the future of the computing infrastructure behind the discovery of the Higgs boson at the Large Hadron Collider.

"We warmly welcome the European and worldwide distributed computing community to Helsinki, a city right at the Baltic sea, which is at its most beautiful in May-June", says Professor Kai Nordlund of the University of Helsinki.

EGI supports 22,000 researchers working in multi-disciplinary science and delivers unprecedented data analysis capabilities to global collaborations. The event will focus on the

current and potential user communities that rely on EGI and will feature training events, demonstrations and a number of technical sessions and workshops. The discussions will help the community to shape the future of EGI. Key themes will be science gateways and workflows, pilot use cases of the EGI Federated Cloud and the development of Virtual Research Environments

The Community Forum will take place at University of Helsinki's main building in the heart of Helsinki. The event will be co-located with the several satellite meetings, such as the EUDAT (European Data Infrastructure) annual event.

"We are honoured to host this very important meeting at the University of Helsinki. Hosting the distributed computing

community is very interesting for us, as our University and CSC have extensive activities in both grid and cloud computing", adds Nordlund.

The EGI Community Forum is hosted by EGI.eu, not-for-profit foundation established to coordinate and manage the EGI federation on behalf of its participants, the University of Helsinki, and CSC – IT Center for Science Ltd.

More information

Information on the event and its programme, keynotes, technical sessions and workshops, will be added to the event's website:
<http://cf2014.egi.eu/>.

Registration to the event starts on January, 13th 2014.