

Inside this  
issue:

**Page 3**  
**Collaborations**  
**Workshop**  
**round up**

**Page 6**  
**NGS**  
**Applications**  
**Corner**

**Page 8**  
**NGS**  
**Sysadmins**  
**Corner**

**Page 9**  
**Collaborations**  
**Corner –**  
**University of**  
**Huddersfield**

**Page 10**  
**NGS User**  
**Case Study**

## Meet the Champions!

One of the single biggest barriers to using the e-tools and e-infrastructures is the knowledge that they exist. To tackle this NGS and SSI have formed a network of Community Champions, scientists who are actively using e-infrastructures in their

research. The project is called Supporting e-Infrastructure Uptake through Community



Champions for Research (SeIUCCR) and it is funded by EPSRC for 2 years.

The champions will disseminate opportunities that are relevant and essential to their colleagues, overcoming the traditional barriers to enable high impact science. The champions will be funded to attend meetings and conferences and tell their colleagues their success stories. It'll also fund researchers with an interest in e-infrastructures to attend its workshops.

SeIUCCR currently has 7 funded champions who are:

- Dr David Wallom (Oxford University) – Energy and Cloud Computing
- Prof Simon Cox (University of Southampton) – Engineering
- Dr Pamela Greenwell (University of Westminster) – Molecular Modelling
- Dr Phil Fowler (Oxford University) – Computational Biophysics
- Dr Rebecca Notman (University of Warwick) – Biological and materials modelling
- Dr Zhongwei Guan (University of Liverpool) – Engineering
- Dr Nicholas Holliman (Durham University) – Physics and visualisation

SeIUCCR will bring together existing networks of advocates in e-infrastructures e.g. SSI's PALS and NGS Campus Champions. The cross-disciplinary nature of the networks will enable the program to develop to include a growing community of multi-disciplinary researchers who can talk the language of application *and* technology.

SeIUCCR will hold a number of events including a summer school (Sept 2011) and will offer up to 20 fully subsidised places on the residential course. Delegates will learn about using e-infrastructures, from local campus resources, through national and international grids to clouds. They will learn about data management and workflows and there will be drop in sessions on fundamentals such as Linux. Importantly, Champions will be there to describe how they have benefited from using the technologies and experts are on hand to solve their real issues.

For more information about SeIUCCR please email [claire.devereux@stfc.ac.uk](mailto:claire.devereux@stfc.ac.uk). Check the NGS website and news email bulletin soon for the launch of the SeIUCCR website and Summer School registration.

## NGS on film

In February UCISA held a Cloud Computing seminar for the UK academic sector. It was billed as "a perfect opportunity to take a look at the innovation and challenges with Cloud Computing and to reflect on the benefits it can bring to your organisation".

The sessions were designed to give a flavour of Cloud Computing possibilities as well as case studies from early adopters and covered topics such as Software as a Service (SaaS) to full Infrastructure as a Service (IaaS).



For those that were unable to attend the meeting, the videos of all the presentations are now available online (<http://www.ucisa.ac.uk/en/groups/ng/Events/2011/Cloud/videos.aspx>) including a presentation from the NGS Technical Director, David Wallom. David presented on the FleSSR project - Flexible Services for the Support of Research (<http://flessr.blogspot.com/>).

David also starred on film at the International Symposium on Grids and Clouds 2011 (ISGC 2011; <http://event.twgrid.org/isgc2011/index.html>). He was interviewed at ISGC by the GridCast team regarding the future of the NGS in his role as Vice President of Community. The video is available from the GridCast blog (<http://gridtalk-project.blogspot.com/2011/03/david-walloms-interview.html>).

## NGS Surgery Meetings

The NGS Surgery meetings (<http://www.ngs.ac.uk/surgery-meetings>) have previously focused on supporting NGS member sites installing Grid interface software on their resources. However recent surgery meetings have widened their remit to discuss topics of interest to Campus Champions (<http://www.ngs.ac.uk/campus-champions>) as well as the Sys Admin community within the NGS.

Suggestions for topics for presentations at NGS surgery meetings are always welcome and they can be emailed to the NGS Helpdesk ([support@grid-support.ac.uk](mailto:support@grid-support.ac.uk)). Details of earlier NGS surgery meetings including links to presentations are available on the NGS Surgery website pages.

If you would like to be sent announcements for surgery meetings then please contact the helpdesk at the above email. The NGS Surgery meetings run every 2nd Wednesday via Access Grid/EVO.

## The fear of being open

**Simon Hettrick from the Software Sustainability Institute (SSI; [www.software.ac.uk](http://www.software.ac.uk)) recaps the successful Collaborations Workshop held in Edinburgh earlier this year.**

Being open has its benefits. If information about your project is accessible to anyone, people will find it easier to get involved. But there's no denying it: opening your work to the scrutiny of the outside world is scary. And this is something I can confirm, because I recently published every slide, every talk and every outcome from the Collaborations Workshop (<http://software.ac.uk/home/cw11/cw11-materials>). And, even scarier, I'm going to publish updates on our progress through the outcomes too.

The Collaborations Workshop, which is run by the Software Sustainability Institute, brings together people from across the research community and gets them to work together on problems faced by researchers. With this community-focussed goal, it would be nonsensical to keep the findings of the workshop to ourselves.

By making everything created at the workshop readily available on our website, we're hoping to include everyone in the collaborative process that is fundamental to the workshop's success. Anyone can view videos of the talks, read the slides and see the outcomes from the discussions. Maybe you agree with our conclusions, maybe you don't... or maybe you think we should focus on different problems altogether. Whatever your view, we want to hear what you think (<http://www.software.ac.uk/about/contact>), because it will help us tailor our work to meet the needs of the research community.

The Collaborations Workshop attracted over 40 attendees who split into small groups and discussed topics like 'How can I get my research used by others?' and 'Should PhD students be taught to program?'. The results of these talks were distilled into outcomes, which are things that the Software Sustainability Institute can do – or ask other organisations to do – to help the research community. We ended up with around 50 outcomes that ranged from lobbying funders to make software sustainability a part of research proposals, to making software training resources more visible on the net.

It's not easy getting people interested in a new idea like the Collaborations Workshop, especially since they need to take a couple of days out of their schedule, a trip to Edinburgh and take part in some intense discussions. And this is another reason for making everything open. We want to show researchers that a flexible workshop with an agenda put together by the attendees on the day of the workshop is not a scary concept, it's one that leads to useful discussions. For that reason alone it's worth being open.

## Getting the Users Story Across

**Gillian Sinclair, NGS Liaison Officer, explains how the NGS is demonstrating the use and benefits of the NGS to a new generation of users.**

Introducing new researchers to the NGS and the services that it offers can sometimes be a difficult task. There are still many researchers who have not yet heard about e-science or the grid and have no idea what they can do with it or what benefits it brings.

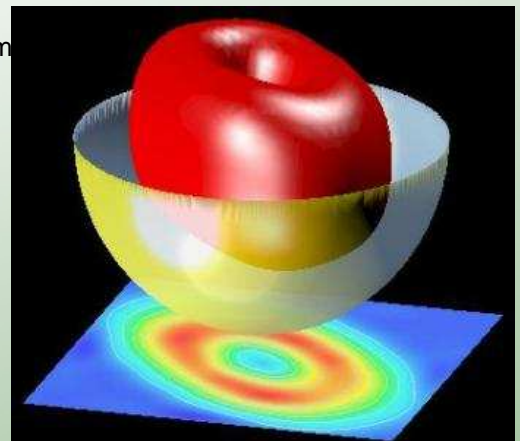
Part of my job as the NGS Liaison Officer is to introduce these researchers to the possibilities that the NGS opens up and to explain how it is relevant to their research. There are many ways to do this - through the website, speaking to researchers at their own research area conferences, through our blog etc. However at the end of the day I have to make the NGS relevant to that particular researcher.

It's impossible to cover every single research area that's out there but here at the NGS we are doing our best to cover as wide a range as possible through our user case studies (<http://www.ngs.ac.uk/case-studies>). We use these case studies to demonstrate how researchers from their field have used and benefited from the NGS.


The case studies are written in conjunction with the researchers and aim to explain (in as much plain English as possible) what they are researching, how they run jobs on the NGS and what benefits it has brought them.

Currently we have 20 user case studies on the NGS website from

- Astronomy databases
- Real-time visualisation of blood flow through the brain
- The effects of defibrillation on the heart
- Geodemographic modelling
- Ion channel simulations and free energy calculations
- Membrane permeation
- Molecule formation from ultracold gases
- Understanding electrical defibrillation of the heart
- High through-put virtual drug screening using the NGS
- Development of a chemical properties database
- Modelling criminal patterns using the NGS



- mRNA analysis using the NGS
- Quantitative genetic analyses on the NGS
- Simulating the universe on the NGS
- Computational modelling of ion channel biophysics
- Simulating carbon nanotubes on the NGS
- Using the NGS to model the climate impact of aircraft emissions
- Exploiting solar power using the NGS
- Quantum Mechanics modelling using the NGS
- Scalable Road Traffic Monitoring using Grid Computing



Whilst the production of user case studies benefits the NGS, they also benefit the researcher featured as quite often the user case studies are picked up by other grid organisations. This was the case with NGS user Cristiano Sabiu who hit the headlines in iSGTW in April (<http://www.isgtw.org/issues/13-april-2011>) with his user case study on his cosmology research.

Cristiano from the University of Portsmouth used NGS resources to study the distribution of galaxies in the universe by creating 2000 mock universes for comparison to actual galaxy distribution in the Sloan Digital Sky Survey (SDSS). By using the NGS he managed to run 20 full scale simulations which required approximately 100,000 cpu hours over the course of a year.

His user case study was picked up by International Science Grid This Week (iSGTW) and appeared as their headline article. iSGTW is a weekly online publication that provides international coverage of distributed computing and the research it supports. They feature articles on all aspects of distributed computing technology, such as grids and clouds from all research areas and are always on the lookout for interesting stories to cover.

The NGS user case studies keep on coming though and the most recent one is on quantum mechanics modelling. Stewart Reed from the University of Leeds used NGS resources to develop new methods of performing accurate computer simulations of quantum mechanical tunnelling (<http://www.ngs.ac.uk/quantum-mechanics-modelling-using-the-ngs>).

Stewart explained how "the NGS provides excellent computing resources with which to perform these calculations. The computational capacity available from the NGS allows larger systems to be studied more accurately than are possible with standard workstations".


## NGS Applications Corner

In each edition of NGS News we will give an introduction to developments from the NGS. In this issue Tom Byrne, a Year In Industry student from STFC Rutherford Appleton Laboratory and Jason Lander from the University of Leeds introduce the new NGS load monitor.

### The NGS Load Monitor

One of the most interesting looking pages on the NGS website is the load monitor (<http://www.ngs.ac.uk/load-monitor>). Clicking the speedometer style gauge on the front page takes you to a listing of all the UK grid sites and their individual loads, along with obligatory colourful bars.

The load monitor has recently been rewritten, and the new load monitor has a few minor aesthetic changes and several major background changes. These changes will make the load monitor more useful for users, and easier to maintain for the website administrators.



The new Load Monitor was designed to be accessible to new Grid users as well as veterans, and so comprises only three columns:

- The hostname of the Grid cluster.
- The number of used jobs slots and the total number of job slots.
- A coloured bar, representing the percentage of job slots occupied on that resource.

The coloured summary bars use a red-amber-green colour scheme to represent resource availability. If it's red, it's busy, and green indicates an idle resource.

### Why is it useful?

The load monitor can be used by users to establish the best host for their jobs. Given that the NGS uses a workload management system (WMS) to distribute jobs to the most suitable resource, one could argue that the Load Monitor is redundant technology.



However, not all sites accessible to NGS users receive jobs from the WMS, instead they must be

accessed directly. For these sites, the load monitor provides an efficient method for checking for usable jobs slots on the entire NGS.

Old-time NGS users will notice that the load monitor no longer references sites, but instead the specific resource host names. While adding a layer of complexity, it increases the tool's usability as users can log straight into the machines using the hostname and their grid credentials.

### How it works

The load monitor is in essence a very simple piece of software. It can be simple because it is built on top of other systems

Host	(Used / Total)	Summary
NGS TOTAL	(42726 / 44573)	 95%
epgr04.ph.bham.ac.uk	(1156 / 1552)	 74%

which do all the hard work, namely, the GOCDDB (Grid Operations Centre's Data Base, where all NGS sites publish site information) and the NGS BDII (where NGS sites publish up to date resource information).

The original load monitor read from a local configuration file which contained the list of sites to report. Unfortunately this file had to be manually updated as sites came and went from the NGS, a design feature which the development team was keen to correct.

The new load monitor instead uses the NGS's single point of truth, the GOCDDB . The GOCDDB provides a list of all UK and Ireland clusters and, using this list, the load monitor queries the NGS BDII servers, which are updated by the individual clusters. This contains the current number of jobs running on each cluster, and the total number of jobs allowed to run on each.

This information is then cached in a local file, and then when a user loads the load monitor, the cached data is retrieved and displayed on the page. If the cached data file is older than 5 minutes, the load monitor will fetch new information from the NGS BDII servers, and update the cache file.

The new load monitor was written by Leeds School of Computing, the back-end is written in PHP by Paul Townsend, and the front-end is Java Script written by Dave Webster. Tom integrated the load monitor into the NGS website with the help of Jason Lander.

## NGS Sysadmins Corner

Every issue we'll be bringing you a handy tip or shortcut from one of the NGS sysadmins. This time it is the turn of Jens Jensen, NGS, STFC Rutherford Appleton Laboratory.

If you ever use the console, here's a tip which will save you a lot of time. Perhaps even your life, if you're of a nervous disposition.

Here's a quick exercise for you: log into a remote machine (Linux or BSD, or whatever.) At this point, you may want to check first that you have all the credentials with you and activated that you need for your daily work - ssh keys, grid proxies, Kerberos tickets.

Now type 'screen'. (If it doesn't give you a welcome screen, complain loudly to your lazy sysadmin.) Hit 'return' and it should give you a console terminal, your first (running your default shell).



The screen tool binds a key, by default control-A, as a control key (which we will abbreviate C-A for now); everything else is passed through to your current terminal. Type C-A c (control-A followed by a lower case 'c') to get your second terminal. Type C-A w to see a list of terminals - the one marked '\*' is your current, the one marked '-' is the previous, and you can switch between those two using C-A C-A. C-A 0 chooses the first, C-A 1 chooses the second, and so on. C-A n and C-A p go to next and previous, respectively.

Try something in the terminals. If you need credentials, you should find that they have all got access to them. If you need a genuine C-A in your input, e.g. with Emacs or Lynx, use C-A a. To exit screen completely, just close all your shells in the terminals (e.g. with C-D for 'normal' shells.)

Now for the great stuff. Type C-A d to detach from screen. You can now close your connection to the machine, but screen will still be there for you, guarding your stuff. If you log in again, possibly from another machine, and type 'screen -r' (for reattach), you should get all your terminals back. You can do work from work, detach, and log in again from home and continue. If you resize your window, screen will send a 'resize' signal to all your terminals.

As if that wasn't enough goodness, if the connection goes down, screen will notice and detach automatically, saving all your terminals. This is much better behaviour than normal ssh logins which will kill everything if the connection goes down (as well as kill people of nervous disposition.)

By now, you should see the power of screen, and we have only just started. You can have screen monitor a terminal for activity (it could be a process waiting for something), or for inactivity (e.g. package updates or compiling.) You can copy from one terminal to another, although with highlighting it doesn't always work so well, you may want to mouse it over instead if you can. You can lock it with C-A x. If you explore, you may well find things I don't know about.

## Collaboration Focus

Each issue we will introduce you to one of the institutions that make up the UK National Grid Service. In this issue the NGS member site of the University of Huddersfield tell Gillian Sinclair why they got involved with the NGS.

**Gillian (GS): Where is the team based?**

**Hudds (H):** The High Performance Computing Resource Centre at the University of Huddersfield is an initiative by the HPC Research Group. The team is based in the Queensgate Campus.

**GS: Who is in the team?**

**H:** Dr Violeta Holmes, Computing & Engineering - research group leader; Dr David Cooke, Applied Sciences - RA Manager; Joanna Radley, Computing Services - manager of Networks & Data centres; Dr Paul Elliott, Applied Sciences; Ibad Kureshi, Administrator – sys admin for HPC resources and NGS Campus Champion; Shuo Liang, Software Administrator, and Dave Andrews, Computing and Engineering - chief technical administrator.



**Back row: Ibad, David**  
**Front row: Joanna, Paul, Violeta**

**GS: How long have you been a member of the NGS?**

**H:** Our researchers first started using the NGS in 2008. These interactions led to our own internal HPC development and from the start we wanted to be a part of the NGS.

**GS: Why did you decide to become a member of the NGS?**

**H:** The NGS forms a forum of discussion for young HPC research groups like ours, to gain guidance on how to develop further. The NGS also sponsors the spirit of collaboration which has helped our researchers immensely. Now at a time when the core sites are rolling back we feel that it is important for institutions that have benefited from the “free resources” to contribute back so that other new sites can enjoy the same.

**GS: What advantages does being a member site bring Huddersfield?**

**H:** Being part of the NGS helps to put us at the cutting edge of research, not just in the field of HPC. It also allows us to put our system through its paces and develop tools based on actual needs from the e-science community. Our presence within the NGS also gives us further opportunities for collaboration.

**GS: What projects are the Huddersfield team participating in that involve the NGS?**


**H:** David Cooke and Paul Martin were part of the Mott VO working on various solid state projects using Castep and DLPOLY on the wider NGS. Paul Elliott has been using Gamess-UK and NWChem to perform gas phase DFT calculations on a range of photophysically active transitional metal complexes. On the service side, we are looking to start work on adopting a campus wide WMS system based on some ideas from David Wallom.

## Case Study

Case studies featuring research carried out by NGS users are regularly featured on the NGS website but in NGS News we will bring one of the case studies directly to you! If you would like your research to be featured on the NGS website and in NGS News as a case study then please contact the NGS Liaison Officer, Gillian Sinclair ([Gillian.sinclair@manchester.ac.uk](mailto:Gillian.sinclair@manchester.ac.uk)).

### Computational modelling of ion channel biophysics

Signals in the vertebrate nervous system move along a neuron as tiny electrical impulses. At the interface, these signals hop between neurons with the aid of a category of specialised proteins called neurotransmitter receptors. These receptors are in the form of ion channels that allow ions to cross the, otherwise impermeable, cell membrane of the neuron.



Ranjit is studying the fundamental principles that drive these neurotransmitter receptor molecules. Understanding the behaviour of these molecules, both in the healthy and diseased states, is critical in designing therapies for a range of neurological disorders including Alzheimer's disease, Parkinson's disease and epilepsy.

Molecular modelling, ligand docking and molecular dynamics simulations allow Ranjit to explore the structure, function and dynamics of these molecules at an atomic level. His research is largely focussed on how the behaviour of these molecules is altered by the binding of small molecules and ions present in its environment. Other minor interests also include how certain short protein fragments could unfold and assemble into fibrils implicated in Alzheimer's disease as well as how water molecules interact with proteins at an atomic level.

Ranjit makes extensive use of the NGS clusters to perform rigorous molecular dynamics simulations and free energy calculations to look at the energetics of physiological processes like ligand and ion binding. Recently, using a range of computational techniques, Ranjit identified the precise location of modulatory ion binding sites in a class of neurotransmitter receptor called ionotropic glutamate receptors. The location of the binding site had puzzled experimentalists for many years without much success but through using NGS resources, Ranjit achieved a positive identification.

Ranjit said "I have been using the NGS for over 4 years, first as a doctoral candidate and now as a

postdoctoral researcher. The NGS and the NGS team have been an immense help to me. Quite a lot of the data reported in my PhD thesis was generated on the NGS clusters. I can honestly say that without the NGS I would probably not have been able to complete my PhD on time”.

Ranjit's supervisor, Dr Phil Biggin, agrees "The NGS has been a fantastic resource throughout the duration of this project. Many of the results would simply not have been possible without the additional computational resource that the NGS provides. It has also afforded us more flexibility in planning which calculations we should give priority to - something that is not always obvious from the outset."

## Updates to NGS Resources

### R installed on the NGS

R version 2.12.1 (R\_2\_12\_1) is available on the NGS RAL Tier 1 affiliate site (lcgce05.gridpp.rl.ac.uk). This is the first NGS UEE compliant application installed on this affiliate site.

R is a data analysis language and environment for statistical computing and graphics that provides linear and non-linear statistical modelling for scientists and researchers. More information regarding R can be found at [http://www.ngs.ac.uk/applications/data\\_analysis/r](http://www.ngs.ac.uk/applications/data_analysis/r).

The installation is done on Scientific Linux 5 (SL5) systems and users are able to run the application via a UI/WMS job submission.



### New resources at the University of Lancaster for NGS users

The NGS site at Lancaster is retiring its old compute cluster and replacing its NGS contribution with access to the new High-End Computing Cluster.

The HEC consists of over 1700 Nehalem cores behind an LSF scheduler, and offers over 3GB of RAM per user job. Access to NGS users is provided with help from the local GridPP team, and can be reached using the NGS WMS.

More information about this cluster and how it can be utilised can be found on the Lancaster member site webpage (<http://www.ngs.ac.uk/lancaster>).



Event	Date	Location	Website	NGS Presence
OGF 32	15 <sup>th</sup> – 18 <sup>th</sup> July 2011	Salt Lake City , USA	<a href="http://www.gridforum.org/OGF32/">http://www.gridforum.org/OGF32/</a>	
EGI Technical Forum	19 <sup>th</sup> – 23 <sup>rd</sup> Sept 2011	Lyon, France	<a href="http://tf2011.egi.eu/">http://tf2011.egi.eu/</a>	NGS staff will be attending.
12th IEEE/ACM International Conference on Grid Computing (Grid2011)	21 <sup>st</sup> – 23 <sup>rd</sup> Sept 2011	Lyon, France (in conjunction with the EGI Technical Forum)	<a href="http://grid2011.mnm-team.org/">http://grid2011.mnm-team.org/</a>	NGS staff will be attending.
UK e-Science All Hands Meeting	27 <sup>th</sup> – 29 <sup>th</sup> Sept 2011	York, UK	<a href="http://www.allhands.org.uk/">http://www.allhands.org.uk/</a>	NGS staff will be attending.
7th IEEE International Conference on e-Science	5 <sup>th</sup> – 8 <sup>th</sup> Dec 2011	Stockholm, Sweden	<a href="http://www.escience2011.org/">http://www.escience2011.org/</a>	

## How to get involved

We hope you have enjoyed this latest edition of NGS News. If you have any NGS-related or grid computing news that you would like to see published in the next edition (September 2011) then please contact Gillian Sinclair ([gillian.sinclair@manchester.ac.uk](mailto:gillian.sinclair@manchester.ac.uk)).

If you are a researcher currently using the NGS or if you have recently used the NGS, we would like to feature your research as a case study for our website and to possibly use it as promotional literature for the NGS. We are also looking for research papers, conference presentations and posters to feature on our website. If you are interested in distributing your research findings to a wide UK and international audience then please contact Gillian at the address above.