

Inside this  
issue:

**Page 3**  
**Updates**  
**from the**  
**NGS**

**Page 4**  
**NGS**  
**Applications**  
**Corner**

**Page 7**  
**Sysadmins**  
**Corner**

**Page 9**  
**Ask Steve!**

**Page 11**  
**User Case**  
**Study**

**Page 13**  
**NGS**  
**Presents...**

## All Change at the NGS

The NGS is undergoing major changes over the next few months. We are strongly encouraging all users to ensure they are aware of these changes and the effect they may have on their usage of the NGS.



At the end of March 2011 approximately 700 free-to-use CPU cores and most SRB and Oracle data-base resources from core

funded NGS sites will be retired from service. This is due to the age and efficiency of the hardware and to allow the NGS to transition to a sustainable platform to better support e-research.

The CPU resources being withdrawn from service represent less than half of core funded sites and only a fraction of the total number of CPU cores that are available to the NGS Virtual Organisation (NGS VO) and its users. Every user with an account on the NGS is automatically a member of the NGS VO and this VO is supported by most of our numerous partner and affiliate sites.

To continue to be able to support early career researchers and pump prime novel projects with a smaller resource pool, all individual user accounts exceeding 2000 CPU hrs allocation will be reset to 2000 CPU hrs on 1st March 2011.

For full details of the announcement and for further details on the free-to-use core NGS resources please see the NGS website (<http://www.ngs.ac.uk>) where all details including a description of resources post March 2011, is available.

If you have any comments or concerns regarding the changes to the NGS, please contact the NGS helpdesk ([support@grid-support.ac.uk](mailto:support@grid-support.ac.uk)).

## The NGS at the University of Huddersfield

The NGS outreach team, in conjunction with local organiser Ibad Kureshi, organised a NGS roadshow event at the University of Huddersfield in January. A selection of speakers took part in this event. The NGS Technical Director, David Wallom from the University of Oxford gave an introduction to the NGS as many people in the audience were not familiar with the NGS.

Following David there were presentations from NGS users including Paul Martin, a researcher at the University of Huddersfield. Paul has been using the NGS for computer modelling of thoria in order to determine its suitability for a next generation nuclear fuel. He explained how he has used DL\_POLY 2 which scales very well on the NGS. Paul in particular praised the good on-line instructions/training, FAQ's, blog and helpdesk/support.

Matt Smith from the University of Liverpool spoke about his use of Abaqus on the NGS to model lattice structures. This was a very interesting talk helped by the sample materials Matt brought along to illustrate his presentation. Matt's presentation prompted plenty of questions about his usage of the NGS and about his research.

David Fergusson from the training team at NeSC rounded off the NGS session by explaining how to get started on the NGS and how to apply for a certificate, run jobs etc. The final presentation was given by Ibad Kureshi who outlined the locally available resources and how these tied in with the resources offered by the NGS. All the presentations from the event are available on the NGS website from the event page (<http://www.ngs.ac.uk/university-of-huddersfield-roadshow>).

### Thank you from the NGS

The NGS would like thank everyone who completed the annual NGS user survey this year. We received our biggest ever response and all the replies will be collated into a report which will be made available via the NGS website.

We would like to congratulate the winners of the survey prize draw who each receive an Amazon voucher -

- Paul Martin, University of Huddersfield
- Krishnamoorthy Arumugam, University of Manchester
- Ranjit Vijayan, University of Oxford

## Updates at the NGS

A round up of updates from various NGS member sites.

### CASTEP update at NGS site STFC RAL

The latest version of CASTEP (5.5.1) is now installed on the NGS at STFC-RAL (ngs.rl.ac.uk and the affiliate ui2.scarf.rl.ac.uk).

CASTEP is a software package which uses density functional theory to provide a good atomic-level description of all manner of materials and molecules. CASTEP can give information about total energies, forces and stresses on an atomic system, as well as calculating optimum geometries, band structures, optical spectra, phonon spectra and much more. It can also perform molecular dynamics simulations.

More details of how to register to use the code and run it on the NGS can be found on the CASTEP page of the NGS website (<http://www.ngs.ac.uk/applications/chemistry/casteputils>).

### MEG Service at NGS site Leeds

ngs.leeds.ac.uk is now providing a MyProxy Enabled GSISsh (MEG) service on: ngs.leeds.ac.uk port 2223 The MyProxy enabled GSISsh service provides command line access to a grid-enabled computer from any ssh client.

To use it, upload a 'proxy' certificate to a myproxy.ngs.ac.uk using a tool such as the Certificate Wizard (<http://www.ngs.ac.uk/tools/certwizard>).

You will need to choose a username and password to protect that certificate.

You can then use the username and password and password to log in to ngs.leeds.ac.uk and ngs.rl.ac.uk. The certificate will be collected 'behind the scenes', used to log you in and made available for submitting jobs to the grid.



### STFC-RAL SCARF Site Provides Testing Queue for all NGS Users

The NGS affiliated cluster SCARF at STFC-RAL has generously provided access to a 'test' queue in response to demand from the NGS Innovation forum (<http://www.ngs.ac.uk/events/ngs-IF10>). All NGS users have access to this 6 hour queue from the NGS UI/WMS service (<http://www.ngs.ac.uk/ui-wms>), for the development of new codes in a large (1000+ CPU) cluster MPI environment. Codes developed on SCARF will work on most other systems on the NGS and MPI codes will run without recompilation for longer runs on the NGS partner site STFC-RAL ngs.rl.ac.uk.



## 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 introduces a way to securely accessing a GUI.

### **NX over GSISsh**

The NGS grid has some of the most advanced scientific simulation programs which can be used to study the outcome of real events, at scales that would take many years on conventional computers.

It seems ironic that in order to use these resources you generally have to interact with the command line, an interface which was superseded by the graphical user interface (GUI) on desktop systems in the 1990s. In this article I'm going to explain one of the reasons why that is, and talk about a method of securely accessing a GUI.

One reason for using the command line is security. Secure shell (SSH) allows users to log into a remote machine securely, with the authentication and connection encrypted. This means that people listening in on the connection will not be able to determine your login information, and so a relatively weak method of authentication such as username and password can be used safely.

SSH supports a feature known as X forwarding. This allows graphical applications to be displayed on a remote machine as if they were local applications, provided both machines are running an X server (the Linux windowing system). This is very useful, but has a few disadvantages:

- High bandwidth requirements, as the X windowing information is sent uncompressed, making this unfeasible for anything other than local high speed networks.
- Client machine must be running an X server, which means the client must be a Linux or UNIX machine, or have an X server installed.
- The forwarded program is only running as long as the SSH session is connected, and when your local machine loses connectivity or is shut down, the remote program will also close down.

### **Enter NX**

NX is an X compression technology. It uses a variety of means to compress the X window traffic, allowing the X windowing protocol to use as little as a hundredth of its usual bandwidth. This makes it usable

across wide area networks. It also allows you to disconnect and reconnect from X sessions, enabling you to leave sessions running on the server. It also has a cross platform NX session viewer, which allows you to connect to your NX sessions from most computers (Windows, Mac and Linux).

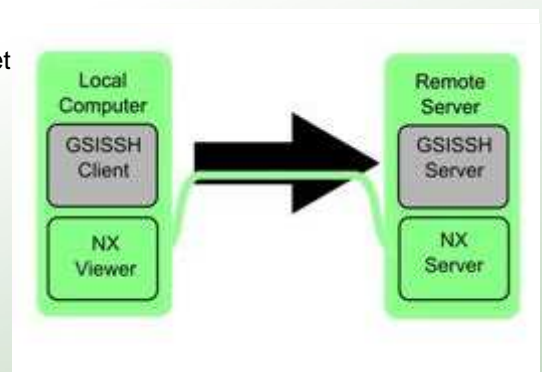
### The Problem

The NGS grid has an extra layer of security on top of SSH, which uses Grid Security Infrastructure to authenticate users via their UK e-Science certificates. This is called GSISSH, and causes a problem for NX.

NX uses their custom implementation of the SSH protocol for authenticating connections to the NX server and setting up a secure communication channel. GSISSH is only widely used in the grid community, and so has not been included in the NX SSH implementation. This means that it would not be possible to connect to a Grid server running a GSISSH server with a NX viewer, as the NX viewer will not support the required authentication methods and will be rejected.

### The Solution

A useful feature of SSH is port tunnelling. This allows you to set up a 'pipe' which will accept data packets from one machine, and spit them out at another machine. The packets travel via the SSH connection, and so have the same level of security as SSH.



This means you can set up a GSISSH connection with the server, and create a tunnel from a port on the local machine to the NX server port on the remote machine. Then, when you point NX viewer at the local end of the tunnel, the NX viewer's data packets will be transmitted across the GSISSH tunnel and be accepted by the remote computer, as they have been authenticated.

### Final Thoughts

While this is a handy trick, it is unlikely to be widely implemented for several reasons.

- Running X sessions requires greater resources than a simple command line login, and the load will be concentrated on the head node of the cluster you are connecting to. This could lead to the head node getting bogged down, rendering the entire cluster unreachable.



- Most grid tools are command line based. Tools like the Globus toolkit, used for submitting jobs to the cluster use a command line interface, and so you do not gain much from having a GUI on the cluster.

However, it has some uses:

- Allowing users who are less familiar with the command line to use a file browser.
- Allowing you to work on the cluster, and then suspend your session and continue working from a different location.

NX over GSISSH has been implemented on the White Rose Grid, a computing resource to users located at York, Leeds and Sheffield universities, but is only available to the university users at York at the moment.

## QMUL joins the NGS

The NGS is pleased to welcome Queen Mary University London (QMUL) as a new affiliate member enabling wider use of their resources through the UK.

QMUL is a Tier 2 site therefore playing an important role in processing and analysing the data produced from the Large Hadron Collider (LHC), as well as being a member of the UK GridPP project.

The university will be contributing their High Throughput Cluster, which is one of the largest sites within the UK in terms of storage and processing power. It is optimised for the analysis of large scale datasets, such as the LHC and currently consists of over 2000 CPU cores interfaced to 300 TB of storage. This will increase to 3000 CPU cores and 1PB of storage in the near future.

Alex Martin, QMUL Grid Project Manager explained "Our team has been providing computing resources through many grid projects, including GridPP and EGI, for almost 10 years, and in that time we have supported many different disciplines both across the UK and further afield. For us joining the NGS, is the next important step in helping to provide a sustainable distributed computing resource for UK researchers, no matter what area they work in".

Chris Walker, Grid cluster manager at QMUL explained how users would benefit at Queen Mary "This agreement will also benefit the university, with the grid team broadening the scope of their work but also giving the researchers here greater access to the NGS, increasing the facilities open to them".

## 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 Jonathan Churchill, NGS, STFC Rutherford Appleton Laboratory.

Have you used the NGS UI/WMS (<http://www.ngs.ac.uk/uiwms>) to submit jobs, and then not been able to retrieve all of your output files, if the job crashed ? Or perhaps you don't know which output files are produced by a program, and just need to look at everything ? Using the "Epilogue" feature in your JDL files you can tar up the complete output directory at the end of a job and return that to the UI machine, to achieve this.

The "Epilogue" is a command that can be run after the main part of the job completes and it can take "EpilogueArguments". The 'trick' in using the epilogue feature is to make sure the Epilogue command always returns an error code of zero. So for example in the JDL we have to use:

```
#!/bin/bash
/bin/tar -czf $1 $2 2>/dev/null
if [ $? -gt 0 ]; then
  echo "Non zero return code from tar. Continuing anyway"
fi
exit 0
```

Rather than:

```
Epilogue = "/bin/tar";
EpilogueArguments = "-cvf outputs.tar.gz *.*";
```



To return the tar gzipped outputs.tar.gz you need to include it in your OutputSandBox list of files e.g.

```
OutputSandbox = {"std.out",
  "std.err",
  "outputs.tar.gz"
};
```

And you have to upload the shell script to the grid resource running the job, so include it in your InputSandBox list e.g. :

```
InputSandbox = {"myjob1.param","myjob1.in","epilogue.sh"};
```

So for example a CASTEP job JDL (see <http://www.ngs.ac.uk/applications/chemistry/castep> for the same without the Epilogue) :

# CASTEP

```
Type = "Job";
JobType = "mpich";
Executable = "/usr/ngs/CASTEP";
Arguments = "CastepTest";
CpuNumber = 16;
StdOutput = "CastepTest.out";
StdError = "CastepTest.err";
InputSandbox = {"CastepTest.param", "CastepTest.cell", "epilogue.sh"};
InputSandboxBaseURI = "gsiftp://ngsui03.ngs.ac.uk:2811/home/ngsxxxx/castep";
OutputSandbox = {"CastepTest.out",
                 "CastepTest.err",
                 "outputs.tar.gz"};
};
OutputSandboxBaseDestURI = "gsiftp://ngsui03.ngs.ac.uk:2811/home/ngsxxxx/castep/outputs";
Epilogue = "epilogue.sh";
EpilogueArguments = "outputs.tar.gz *.*";

#
# Forces to job to run on any site with CASTEP installed.
Requirements = (
    member("NGS-UEE-CASTEP", other.GlueHostApplicationSoftwareRunTimeEnvironment)
);
Rank = other.GlueCEStateFreeCPUs;
ShallowRetryCount = -1;
```

The same approach can be taken if your own programs have more than 20 output files, as this is the NGS WMS OutputSandBox file limit.

Of course all kinds of things are possible with an Epilogue script, not just taring up all the output files. It can be used to mail you when the job is completed, or run a program to extract some meta data from your run and send it to a database somewhere.

Users of decommissioned NGS RMCS service will recognise this ability as where the "RCommands" were run, or in the gLite world where AMGA (ARDA metadata Catalogue Project <http://amga.web.cern.ch/amga/index.html>) commands could be run, to upload metadata to a central database.

Attendees at the Innovation forum noted that the heavily used machines on the NGS made developing new code difficult. By providing this fast turn around queue on RAL-SCARF, it is hoped that code development can be sped up for NGS users.

The SCARF cluster is provided by STFC for their researchers and collaborators. For example at ISIS, Diamond and the Central Laser Facility. In addition several, exclusive access parts, parts of SCARF are run on behalf of other research groups and institutions in a shared service facility.

## Got a question about software? Ask Steve!

**Simon Hettrick from the Software Sustainability Institute introduces their new service.**

In 2010, a crack developer was asked to join the Software Sustainability Institute (SSI; [www.software.ac.uk](http://www.software.ac.uk)).

This man promptly set up as the Institute's software architect. Today, still wanted by a number of projects, Steve survives as a developer of fortune. If you have a problem, if no one else can help, and if you can email him, maybe you can Ask Steve!

At the SSI, whenever we have a software problem we simply ask Steve. He's our in-house software architect and all-round guru of code. Then we got to thinking: it's selfish to keep such a valuable resource to ourselves, we should make Steve's knowledge available to everyone. And that's when the idea for the Ask Steve! blog (<http://asksteve.software.ac.uk/>) was born.

So what is Ask Steve? The idea is that anyone can email Ask Steve! with a software trouble and query. Each week or so, Steve will work on a problem and post his answer to the blog. People can comment, try out the solution or simply get back to Steve with another question. Steve will sort through the questions he is posed and answer the ones that trouble the most people.

Next time you have a software problem, visit the Ask Steve! blog or email Ask Steve! ([asksteve@software.ac.uk](mailto:asksteve@software.ac.uk)).



## Publicity for our NGS Users

**As well as promoting the NGS as a service, the NGS Liaison Officer also promotes the research performed on the NGS by our users. In this article the NGS Liaison Officer, Gillian Sinclair explains how users research gets the publicity it deserves.**

There have been some new NGS user case studies (<http://www.ngs.ac.uk/case-studies>) added to the website over the last 3 months and there are more to come. The user case studies are a collaboration between the user and myself as the NGS Liaison Officer. I work with the user to produce a short case study outlining their research, how they use the NGS and, rather importantly, the benefits that the NGS has brought them.

Most of the user case studies come about through users volunteering to write a case study, mainly through our annual user survey. For example this year 31 users have volunteered to produce a user case study and I will be contacting them in turn over the next few months.

I hope the increase has come about through people realising the benefits of advertising their research to the wider community. The most enthusiastic communicators in my experience for the NGS, tend to be PhD students which bodes well for the future of science communication.

The user case studies don't just stay on our website. With over 2000 hits since they were placed on our website, they have been picked up by other dissemination teams UK and Europe wide. Quite a few of our case studies have been picked up by iSGTW (<http://www.isgtw.org/>) which has 7300 subscribers and many thousands more unique visitors to their website. EGI (<http://www.egi.eu/>) have also become interested in our case studies and are looking to produce something similar for their own project.

And it's not just within e-science / grid / e-infrastructure organisations that the case studies are picked up. NGS user research features frequently in Scientific Computing World (<http://www.scientific-computing.com/>). I am often asked by editors if I know of anyone researching X or Y and, thanks in no small part to the user case studies and the NGS Communities service (<http://www.ngs.ac.uk/Communities>), I can usually find a researcher for them to interview. SCW is free to read online and you can also subscribe to a free printed copy.

The latest NGS user case study has recently been put on our website - Simulating carbon nanotubules on the NGS by Rebeca Garcia Fandino at the University of Oxford (<http://www.ngs.ac.uk/simulating-carbon-nanotubules-on-the-ngs>). Watch this space for more new case studies coming soon!

## 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)).


### Simulating the universe on the NGS

Cristiano Sabiu is a recent STFC funded PhD graduate from the University of Portsmouth where he studied the distribution of galaxies in the universe under the supervision of Prof. Robert Nichol. The distribution of galaxies can tell us a great deal about the underlying physics which govern the dynamics and evolution of the universe.

There are currently two unsolved issues in cosmology. It is known from recent observations that the universe is not only expanding but is doing so at an accelerated rate, but what is driving this? This unknown force is usually called dark energy. The second issue is that it is known from observations that an extra source of gravity must exist in the universe. It is thought that there are particles which do not emit light and only interact through the gravitational force, this is usually called dark matter.

It is thought that galaxies form where there is a concentration of dark matter but the exact relationship between the dark matter and galaxies is hotly debated. The relationship is usually characterised by a very simple function known as 'bias'. Cristiano used an alternative to the standard bias, known as the Halo Occupation Distribution (HOD). The HOD is a method for placing galaxies within a Dark Matter distribution. Although the HOD has many advantages over the bias prescription, its exact form is not well understood.

His research focused on finding the exact form of the HOD, mathematically this involves solving for the unknown variables in an equation but in reality he is exploring how the galaxies we see are connected to the invisible dark matter. This required running many large scale N-body simulations of dark matter. From these simulations he created a multiverse of almost 2000 mock universes which explored multiple HOD models. These were then compared to the actual galaxy distribution as observed by the Sloan Digital Sky Survey (SDSS) allowing the placing of tight constraints on the model parameters.





The simulations were run using the publicly available GADGET2 code. GADGET2 is an N-body integration routine which can run cosmological simulations on massively parallel computers with distributed memory. GADGET2 uses an explicit communication model that is implemented with the standardised MPI communication interface. The code can be run on essentially all supercomputer systems presently in use, including clusters of workstations or individual PCs.

Cristiano simulated the large scale structure using over 100 million dark matter particles and, as his models were generally very spatially homogeneous, the code showed excellent scalability. He completed a series of cosmological simulations, using a 128 CPU configuration on the NGS-RAL facility and successfully ran 20 full scale simulations, requiring ~100,000 CPU hours over a 12 month time period.

Cristiano said "GADGET2 was installed for me on the NGS and optimised for their system. It is now available to all NGS users so if you fancy creating your very own Universe why not give it a go!" adding that "Without the NGS my PhD project could have taken 10 years!".

## Call for proposals for commissioned technical materials

The ETF network (<http://www.escience-etf.ac.uk/>) is inviting all e-Science or related projects to bid for funds to enable the capture of technical detail which may otherwise be lost at the end of a project.

The ETF network is extending this invitation to all past, current and proposed e-Science or related projects to bid for funds to enable the capture of technical detail which may otherwise be lost at the end of a project.

These funded dissemination activities may take the form of a report/white-paper describing specific technical aspects of the project or the development, maintenance or extension of a specific HOWTO document. All of these outputs will be published on the ETF and NeSC (<http://www.nesc.ac.uk>) information network websites.

For more information see the news article (<http://www.ngs.ac.uk/news/call-for-proposals-for-commissioned-technical-materials>) on the NGS website.

## NGS presents...

Each issue we introduce a member of NGS staff who will explain their role within the organisation. In this edition we are delighted to introduce Dr Steve Brewer from EGI (European Grid Initiative).

I work as Chief Community Officer at EGI.eu (<http://www.egi.eu>) based in Amsterdam in the Netherlands. My job is to be the voice of the user in terms of steering the evolution of the EGI e-Infrastructure. This sounds rather grand but in reality involves talking to a lot of researchers in many different fields and understanding their research needs with respect to using our e-Infrastructure. In order to dig into these needs and requirements we have a suite of services to support communities and a small team of experts to investigate the researcher's needs and feed the results of these investigations on to the relevant solution providers.



The linchpin of these services is our Requirements Tracking system (<http://www.egi.eu/user-support/>). Users and Communities can submit requirements into this system through a variety of channels ranging from email to direct access to the system. The advantage of our system is its transparency, anybody can connect to it and see what is happening to requirements as they are processed. This means that solution providers from outside EGI can directly view the emerging needs of the community and insert comments and feedback on the process.

We are integrating existing services and assisting research communities and NGIs to embed them within their own portals and gateways as personalised resources. The services currently include an applications database, various VO monitoring tools aimed at research groups and a training calendar and database.

We are also active in cultivating and nurturing research communities. Research communities must be of a sufficient size to give the community a voice which is where the EGI User Community Support Team (UCST) comes in. We can help communities by providing tools and services as well as events and workshops but ultimately they have to find that voice from within themselves.

Traditionally, the dominant voice in this area has come from the team based around the LHC experiments at CERN, and this will continue to be the case for the years ahead. However, new voices are beginning to emerge in other equally interesting and challenging areas. These include Structural Biology, Hydrometeorology and, increasingly, the Arts and Humanities. Techniques such as 3-dimensional visualisation and geo-spatial mapping can be utilised by many disciplines so the sharing and sign-posting of knowledge has never been more imperative.

Event	Date	Location	Website	NGS Presence
Collaborations Workshop	3 <sup>rd</sup> – 4 <sup>th</sup> March 2011	Edinburgh, UK	<a href="http://www.software.ac.uk/events/2011-03-03-collaborations-workshop">http://www.software.ac.uk/events/2011-03-03-collaborations-workshop</a>	NGS posters and NGS staff will be in attendance
JISC annual conference	14 <sup>th</sup> – 15 <sup>th</sup> March 2011	Liverpool, UK	<a href="http://www.jisc.ac.uk/events">http://www.jisc.ac.uk/events</a>	NGS exhibition stand and NGS session
ISGC 2011 & OGF 31	19 <sup>th</sup> – 25 <sup>th</sup> March 2011	Taipei, Taiwan	<a href="http://event.twgrid.org/isgc2011/">http://event.twgrid.org/isgc2011/</a>	
EGI User Forum	11 <sup>th</sup> – 15 <sup>th</sup> April 2011	Vilnius, Lithuania	<a href="http://uf2011.egi.eu/">http://uf2011.egi.eu/</a>	NGS exhibition stand as part of the NGI
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>	

## 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 (June 2011) then please contact Gillian Sinclair (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.