

The UK's Grid Operations Support Centre and the National Grid Service.

Core Components of the UK's e-Infrastructure

Neil Geddes, Stephen Pickles, Andrew Richards

On behalf of the
Grid Operations Support Centre

Abstract

The National Grid Service has now (as of September 2005) been operational for 18 months and in full production for a year. The service supports over 250 academic users covering a wide range of disciplines, from Astronomy to Social Science. The Grid Operations Support Centre, responsible for central administration of the NGS operations and core services, has also been operational for a year. In this paper we review the first full year of operations of the NGS and GOSC, reflect on the lessons learned and the current state of development grid middleware. Current plans for the NGS and GOSC will also be reviewed.

1. Introduction

The full production National Grid Service (NGS) has now been operating for a year. In that time the number of registered users has grown to over 200 and the core NGS nodes typically operate at 50-80% loading. Over the last year the NGS has grown beyond the original "core nodes" with the membership of three new sites. A formal process for addition of resource providers as full partners or affiliates of the NGS has been developed. The Grid Operations Support Centre (GOSC) and NGS have adopted a conservative approach to the deployment and support of new services, focusing wherever possible on well engineered and supported (Open Source) software which supports the strategic goal of developing into a key core component of the emerging UK e-Infrastructure. The GOSC currently operates the UK's National e-Science Certificate Authority, a central helpdesk, and central project (Virtual Organisation) administration and resource brokering for the NGS.




Figure 1. Growth in registered NGS users to April 2005.




Figure 2. An example of NGS load monitoring

2. Current Status

2.1 Use of the NGS

The NGS offers a range of services focused on data access and computation, however, at present most users make use of a very limited range of computational (i.e. access to CPU time) or data management (principally SRB and Oracle) services. More sophisticated uses such as service integration and data and service aggregation are only recently being investigated as user communities gain confidence that the NGS interfaces will continue to be supported for and provide access to useful (scale of) resources. This in part also reflects the immaturity of the available supporting software, and in part the significant community based developments and education that are required to take full advantage of the opportunities offered. Sustained support for a stable infrastructure will be key to encouraging community initiatives to tackle these latter points.

The approval of projects to use the NGS resources has continued steadily throughout the first year, with allocations ranging from a few to several tens of thousands of CPU hours and

from a few GB to TB of data storage. The subject of projects ranges from astronomy to sociology, with most of the users in the first six months coming from the EPSRC and BBSRC communities. More recently we have seen the emergence of more community oriented projects developing standard “portal” access to the NGS services, bringing these services to a wider audience with only limited interest in learning about “grid computing”.

In parallel with the increasing user base, several new sites have become partners in the NGS, offering a range of services to users. The GOSC now administers a documented and tested procedure for joining or affiliating to the NGS and is now looking for more sites who wish to join. Partner sites benefit from direct access to the shared experience developed so far and direct input to the future development of the NGS.

2.2 Core Sites

The NGS is based around four core nodes, funded by the JISC and CCLRC, and hosted by Leeds University, Manchester University, Oxford University and the Rutherford Appleton Laboratory. Two of these nodes offer “compute” services based upon 64 dual (Intel 3GHz) Linux clusters with a fast Myrinet interconnect. The other two “data nodes” are based on smaller, but compatible, 20 dual CPU clusters with 18TB of attached disk and full access to an Oracle 9i clustered database and Oracle Application Server.

In addition to the core sites, the NGS also provides access to the UK’s HPC facilities, CSAR and HPCx. However, time on these machines must first be obtained through the existing academic and commercial mechanisms.

2.3 New Partners

New partners join the NGS through a well defined process. In summary:

1. Define level of service commitments through a Service Level Definition
2. Agree to and adopt NGS acceptable use and security policies.
3. Run compatible middleware as defined by NGS Minimum Software Stack and verified by compliance test suite.
3. Support monitoring and accounting

The compatible middleware defines the NGS interface for user authentication, data transfer and job submission. This interface is OGSA compliant and the NGS core nodes implement it

using the Globus software distributed with the Virtual Data Toolkit. Further, optional, interfaces then provide for interactive access and distributed data storage via the Storage Resource Broker.

Two levels of membership are supported:

Affiliation or simple connection to the NGS by supporting the common interfaces.

Partnership, where in addition to supporting the NGS interfaces, “significant” resources are made available to all NGS users including the ability to compile code for computational resources. In addition, partner sites recognise additional Certificate Authorities with which UK e-Science programme has reciprocal agreements.

Over the last year, three new full partners have joined the NGS; Bristol University, Cardiff University and Lancaster University. These partners offer access to further computational resources which complement those available at the core sites:

Bristol – provides access to a local condor cluster, initially including 20 2.3GHz Athlon processors arranged in 10 dual CPU nodes.

Cardiff – provides 1000 hours per week on a SGI Origin system comprising 4 dual CPU Origin 300 servers with a Myrinet™ interconnect.

Lancaster – provide access to 8 Sun Blade 1000 execution nodes, each with dual UltraSPARC IIICu processors connected to the NGS via a Dell 1750 head node.

Each of these resources has been committed to the NGS for a period of at least 12 months. Further details of these resources can be found on the NGS web site: www.ngs.ac.uk.

The heterogeneity introduced by these new services has allowed us to develop experience in connecting an increasingly wide range of resources to the NGS; experience which will be invaluable in the future. Making full and effective use of this wide range of architectures may present challenges for users. Providing a common, albeit basic, interface for authentication, authorisation and access is the first step towards supporting more sophisticated usage across such a variety of resources. User experience and usage patterns over the coming year will clearly help to inform future investments in the core nodes and partner sites.

3. User Confusion and The Road Ahead

A principal goal of the GOSC is to deploy and support a Web Services (WS) based grid infrastructure to support UK research. This goal has proved significantly more challenging than originally hoped. It took several years to develop the stable GT2 (Globus Toolkit Version 2) based middleware to a state where it is robust enough to be deployed in a production environment such as the NGS. Re-implementing this knowledge and these developments as robust web services has not proved simple. At the time of the original NGS deployment it appeared that the Global Grid Forum's Open Grid Services Architecture (OGSA) would define the framework and the Open Grid Services Infrastructure (OGSI) would provide the infrastructure for the next logical step. OGSI, however, did not receive widespread support and was replaced by a commitment to the Web Services Resource Framework (WSRF), the standardisation of which is only now nearing completion in OASIS. The rapid rise and fall of OGSI, and the dubious quality of the early OGSI implementations, made existing GT2-based projects, including NGS, somewhat wary of rushing into replacing their existing infrastructure.

In parallel, many WS standards are themselves emerging much more slowly than originally hoped. Many important components are not fully specified (or tested) or are the focus of ongoing discussion in the commercial world. Upheaval around the OGSI migration from "Grid Services" to WSRF also delayed coherent development as has uncertainty about the security (authentication and authorisation) models that will be widely supported in the future. In the UK, for example, the decision by JISC to promote the use of Shibboleth as the authentication system for its nationally provided resources has added to the confusion. Although a sound authentication model, Shibboleth does not yet support or interoperate the dynamic authorisation structures required to support virtual organisations (projects) required in a grid environment. Recent initiatives in the US and UK have only just begun to address this latter issue.

3.1 Strategy for the Future

OGSA remains important to the future development of the NGS bring, as it does, the promise of interoperability. OGSA attempts to understand what are the fundamental

capabilities needed to build grids, and how should these be factored into services? An organising principle such as OGSA is required to reduce massive duplication of effort, and focus on improving capabilities, usability, and robustness. After several years work, OGSA is only beginning to deliver and work on first normative spec (Basic Execution Services) is just starting. However, without OGSA and the framework it provides, we would simply have to invent (yet) another similar activity. A key requirements for the future success of OGSA will be its contribution to and adoption by the major production grid deployments around the world and the commercial IT world.

There are, however, encouraging signs. The Job Submission Description Language standard from the Global Grid Forum looks is about to be adopted by several major grid development and production projects. A common interface to storage, the Storage Resource Manager (SRM) has been agreed across a large number of grid projects (though only a limited set of large scale implementations of this standard exist), and a common information schema, the GLUE schema, is on common use around the world. Each of these activities has taken several years to reach agreement as the various developments have confronted the realities of large scale deployment and use.

A key component of our future success in grid computing will be the ability of the nascent "middleware hardening" activities such as the UK's Open Middleware Infrastructure Initiative, to take these emerging standards and early implementations and turn them into robust and user friendly implementations. The world does not need yet another job submission interface, it needs a robust implementation of the agreed and tested open standards!.

All of the above confusion has made defining a technology roadmap for the GOSC and NGS challenging. It has also made users and potential users unwilling to invest significant effort in developing new higher level services while so much uncertainty surrounded the evolution of the low level infrastructure.

The GOSC has agreed a strategic framework, aligned with related international e-infrastructure projects, to provide the basis for future planning. The main components of this framework are:

1. The NGS authentication and authorisation model is based on GSI and in particular the delegation possible in this model is important.

2. The GOSC will fully support only authenticated and auditable (grid/web) services.
3. Compatibility with European e-infrastructure, EGEE, should determine the strategic direction taken by the GOSC.

In addition clearer plans have emerged around the UK's commitment to Shibboleth and the requirements for interoperability with National e-science authorisation infrastructures. In parallel the requirements of the NGS user communities have developed further. Within these guidelines the GOSC has developed plans for the services and infrastructure to be supported through to the end of 2006 and beyond.

4. Conclusions

The services supported by the GOSC will continue to form the basis of a sustainable e-infrastructure, providing the authentication, authorisation, collaboration and resource sharing tools to meet the requirements of research collaborations and access to large facilities national and international facilities and resources. The NGS cares about:

- alignment/compatibility with leading international Grid efforts
- special requirements of UK e-Science community
- easy migration/upgrade paths
- growth: minimising cost of membership for partners based on standards or standards-track specifications
- proven robustness/reliability
- usability: but this is too big a problem for us to solve alone

Significant changes to NGS Minimum Software Stack will require approval by NGS Management Board on conservative time scales. However, NGS and GOSC will endeavour to work with all users to meet their current and future needs as the principal deliverable for the NGS will always be to help deliver effective and reliable services to support UK research.