



White Rose Grid Node on the UK National Grid Service

Service Level Description

Version Control

Version	Date	Author	Changes
0.1	20 August 2007	J.G.Schmidt (j.g.schmidt@leeds.ac.uk)	Draft
0.2	3 September 2007	J.G.Schmidt	Draft; date correction in the footer
0.3	28 November 2007	J.G.Schmidt	Updated to include recently installed software.
0.4	30 November 2007	J.G.Schmidt	Added software (HDF5, MPIg, R); minor corrections in text.

Aim

To make available resources and services as described below via the UK National Grid Service (NGS) to UK academia.

Applies to

The White Rose Grid NGS node (ngs.leeds.ac.uk) based on a ClusterVision high performance Beowulf-type compute cluster hosted at the University of Leeds, which consists of:

- Master/head node with 4 dual-core 2.6GHz AMD Opteron 64-bit processors, 32GB RAM
- 48 nodes with 2 dual-core 2.6GHz AMD Opteron 64-bit processors, 8GB RAM
- 8 nodes with 4 dual-core 2.6GHz AMD Opteron 64-bit processors, 32GB RAM
- Gigabit network interconnect
- Myrinet high-speed message passing interconnect
- ClearSpeed CSX600 card installed on 1 node
- 3 AMD Opteron redundant storage 2U storage servers
- Two 12TB Infortrend SATA/Fibre RAID units
- One 6TB Infortrend SATA/Fibre RAID unit.

The head node acts as a front-end to compute nodes which are used for executing jobs submitted via Globus Toolkit, the appropriate job manager, and the PBSPro scheduling system. The three RAID units provide the user filestore.

Service Description

The National Grid Service is a collection of nodes, and each node specifies through its own SLD (e.g. this document) the resources (hardware, software and support resources) it is making available to the UK academic community through a middleware/grid interface.

Service Provider

Information Systems Services (ISS) jointly with the White Rose Grid e-Science Centre (part of the School of Computing) at the University of Leeds operate the cluster and support its users as part of the White Rose Grid consortium.



Service Users

All users of the service are required to accept the *Regulations for Use of the NGS* stated at:

<http://www.ngs.ac.uk/NGS-tacu.html>

Inclusions

1. The following grid middleware and related software for the support of the UK research community is installed:
 - a. Pre-WS components of the Globus Toolkit version 4.0.3 from the Virtual Data Toolkit (VDT 1.6.1) release, including GridFTP, GSIssh and GRIS reporting to the central BDII.
 - b. PBSpro (Portable Batch System)
 - c. Grid-mapfile populated from the central NGS VOMS server
 - d. RUS accounting software
 - e. Ganglia
 - f. SRB 3.4.2 client and server/vault
 - g. Scripts supporting the Uniform Execution Environment (UEE)
2. The following software for developing and running programs is available:
 - a. Intel Fortran compiler (Category 1 support)
 - b. Intel C/C++ compiler (Category 1 support)
 - c. Intel Math Kernel Library (Category 2 support)
 - d. PGI compilers (Category 2 support)
 - e. High Performance Fortran (HPF) (Category 2 support)
 - f. MPICH-MX message passing library (Category 2 support)
 - g. Java - Various versions of Sun's Java development environment
 - h. GNU compiler suite (including gcc and gfortran)
3. The following applications packages are available:
 - a. Amber (licence required; category 3 support)
 - b. BioPerl (Category 3 support)
 - c. FFTW (Category 3 support)
 - d. GotoBLAS (Category 3 support)
 - e. GROMACS (Category 3 support)
 - f. GSL (Category 3 support)
 - g. HDF5 (Category 3 support)
 - h. LAMMPS (Category 3 support)
 - i. MPIBLACS (Category 3 support)
 - j. MPIg (Category 3 support)
 - k. NAMD (Category 3 support)
 - l. NETCDF (Category 3 support)
 - m. R (Category 3 support)
 - n. ScaLAPACK (Category 3 support)
 - o. SIESTA (licence required; category 3 support)
 - p. TFBS (Category 3 support)
4. Digital certificates management: accept certificates issued by the UK e-Science Certificate Authority and those CAs with which the UK e-Science Core Programme has agreements. The Certificate Revocation Lists are updated on a regular basis.



THE WHITE ROSE GRID

e-Science Centre

5. Information: *User Guide* for users of all the NGS systems is available from <http://www.ngs.ac.uk>, and standard Linux on-line HELP information is available on the system. Further information for this node is available at: <http://docs.ngs.leeds.ac.uk>
6. Provision of usernames: usernames from ngs0001 to ngs1000 will be provided.
7. Provision of disk space for users' files.
8. Provision of a 1TB vault for SRB.
9. Nagios is used for tracking faults and email alerts to system administrators.
10. Usage accounting records are uploaded to the NGS RUS service.
11. Batch job control: a job queuing system is operated on the ClusterVision cluster, and the successful operation of the system is monitored closely.
12. Specialist advice is available to research users who are assumed to be familiar with UNIX and general programming techniques.
13. Systems support - The system software products are supplied under licences specifying conditions on their use and providing a fault reporting and correction service through the ClusterVision Support Centre.

Exclusions

1. Turnaround time cannot be guaranteed, as the system may be heavily utilised and incoming jobs may be queuing waiting for existing jobs to finish.
2. No backups of users' files are carried out.

Service Level

Quality

1. Ganglia (<http://ganglia.ngs.rl.ac.uk/>) results and Inca (<http://inca2.ngs.ac.uk>) results are published regularly and monitored to identify problems.
2. The NGS machines run the Red Hat Linux based ClusterVisionOS operating systems supplied by ClusterVision and versions are normally kept as up-to-date as possible. The Red Hat licence agreement includes patches and updates.
3. Systems software problems are reported to the ClusterVision Support Centre; hardware maintenance contract is in effect on this equipment.
4. Compilers are supported through the relevant licence agreements with Intel and Portland Group.

Availability

5. The NGS machines will be available at all times subject to: essential planned maintenance to hardware or software; the *At Risk* period; and unplanned stoppages and failures.
6. Unplanned stoppages and failures out of hours will be dealt with next working day; hardware support reflects the contract with ClusterVision.
7. Two weeks' notice is usually given to users (via the NGS web site or the appropriate emailing list) about scheduled maintenance.
8. The designated *At Risk* period for Leeds cluster is between 8:00 and 11:00 on Tuesday mornings. This may be used for low risk maintenance work. Note that the NGS service is also *At Risk* on Tuesday mornings due to JANET network being at risk periods.
9. Inca node monitoring results are displayed via NGS Web site at <http://inca2.ngs.ac.uk> ; email alerts are sent if Inca tests fail.
10. The reliability of the service is monitored by the Nagios system which alerts Leeds system administrators if the node or the service goes down.
11. The reliability of the service is monitored and the availability is presented by the Ganglia system (<http://ganglia.ngs.rl.ac.uk/>).



12. The serviceability targets agreed with ClusterVision are: on-site for the head node, storage units and RAID disks. The slave (compute) nodes are posted to ClusterVision for exchange.
13. Progress meetings, regular email communications and discussions are held between the NGS Operational Team and ClusterVision where the serviceability of the NGS systems and their performance are reviewed.

Filestore

14. Users' home directories and the SRB vault (when installed) are daily mirrored for the purpose of recovering lost files in the event of system failure. No off-site storage of the data is currently being made, and the data are kept for a very short and limited period of time.
15. Users' GASS-CACHE files that have not been accessed for 1 month might be deleted.
16. Scratch disk space for individual users will be available.

Compliance

1. When serious problems occur, the escalation procedure is invoked (see **Escalation Mechanisms** below).
2. Availability statistics are published via Ganglia (see above) and Inca (see above).

Operational Framework

1. User problems and queries should be reported to the NGS Helpdesk (<http://www.grid-support.ac.uk/>) either by phoning 01235 446822 or emailing support@grid-support.ac.uk. The NGS Helpdesk will forward details of queries to the Leeds personnel for Leeds node specific issues.
2. Regional users may report their queries directly to the ISS Helpdesk (email: helpdesk@leeds.ac.uk or phone: 0113 343 3333).
3. ISS operations staff and ISS systems staff at the University of Leeds monitor the availability of this service at periodic intervals during their respective working days, and if required report faults and issues to ClusterVision.

Change Control

New releases of system software are generally agreed in advance with other sites at the NGS Operation Team meetings. Users are informed prior to upgrading taking place via the appropriate NGS emailing list.

Support Category Definitions for Software

ISS at the University of Leeds operates a three level scheme for software support. The levels are defined as follows:

- **Recommended Software (Category 1)**
This is the highest support category and is reserved for software of the highest quality and of wide applicability. In-depth knowledge of the software is available in ISS at the University of Leeds to provide a high level of technical support. Advice is provided on basic and advanced use. Problems will be reported to suppliers, their resolution monitored, and fixes applied if supplied. In addition, training courses may be offered for some products.
- **Supported Software (Category 2)**
The majority of software made available falls into this category. Technical support is provided by ISS at the University of Leeds although the level of support varies according to appropriate staff expertise. Problems will be reported to suppliers, their resolution monitored, and fixes applied if supplied.
- **Provided Software (Category 3)**

Software in this category is provided without technical support. This category is used for a variety of reasons. Some software of a highly specialised nature, for which little or no expertise is available within ISS at the University of Leeds, falls into this category. Another reason for using this category may be that technical support is not provided by the software supplier. Although technical support may not be forthcoming, product maintenance may well be provided.

Problem Severity Definitions

A **priority level** is assigned when a problem is reported to the University of Leeds Helpdesk according to its impact:

- **Priority 1**
Severe impact affecting most users across the campus or regionally/nationally
- **Priority 2**
Major impact affecting many users e.g. a whole department or area
- **Priority 3**
Impact confined to a subset of users
- **Priority 4**
Impact confined to just an individual

Escalation Mechanisms

Problems causing the unavailability of the NGS cluster at the University of Leeds would be categorised as either Priority 1 or Priority 2 and be subject to the following escalation procedure:

1. As soon as the problem is properly characterised by ISS staff, following its observation locally, or a report by the NGS Helpdesk, or the Leeds ISS Helpdesk the system administrator will log the incident into its fault tracking system.
2. Support staff for the service are called in, and for any problem requiring contact with a supplier or maintenance company, this is made as soon as possible.
3. If the problem remains by an hour after the identification, the ISS Group Leader responsible for the service is informed.
4. If the problem persists for half a working day, the ISS Service Manager responsible for the service, the Director and Customer Service manager of ISS are informed.
5. If the problem persists between half and one working day, a news item detailing the problem and any possible workarounds will be created and emailed to NGS users.
6. If the Service Manager responsible for the service concerned is dissatisfied with a maintenance company's (ClusterVision) response to a particular problem, the NGS Executive Director is informed in order to escalate this matter in accordance with the terms of the appropriate agreement.

Less serious problems have a lower priority, but are entered by the system administrator into the ISS fault-reporting system if they are not resolved within 3 working hours. The ISS fault reporting system tracks them and subjects them to different levels of escalation.

Definitions

1. *Amber* - Molecular Dynamics Application (licence required)
2. *BioPerl* - Perl tools for bioinformatics, genomics and life science research
3. *BDII* - Berkely Database Information Index acts as a cache, storing information about the Grid status in its database. It collects information by querying the Grid Index Information Service (GIIS) at each site.
4. *compute nodes* - the main cluster machines which run work through the PBSpro workload scheduling system



THE WHITE ROSE GRID e-Science Centre

5. *FFTW* - collection of fast C routines for computing the Discrete Fourier Transform in one or more dimensions
6. *Ganglia* – a Web-based distributed monitoring and execution system for high-performance computing systems such as clusters and Grids
7. *Ganglia Web site for federated NGS* – <http://ganglia.ngs.rl.ac.uk>
8. *GASS* – *Global Access to Secondary Storage*: libraries and utilities for file I/O to the Globus environment
9. *GASS-CACHE* – local file cache in users' home directories used by GASS
10. *Globus Toolkit* – a toolkit providing a set of services to enable grid applications and the underlying grid infrastructure to inter-operate
11. *GotoBLAS* - Basic Linear Algebra library
12. *GridFTP* – an FTP service with extensions to meet requirements of grid environments
13. *GROMACS*- molecular dynamics simulation package
14. *GSIssh* – a version of ssh (secure shell) that uses the Globus Security Infrastructure (GSI) for authentication with X509 certificates
15. *GSL* – GNU scientific library
16. *HDF5*- technology to manage extremely large and complex data collections
17. *head node* - the front-end machine for job submission and routine tasks
18. *LAMMPS* - molecular dynamics simulator
19. *MPIBLACS* - a linear algebra oriented message passing interface
20. *MPIg* -MPIg is a version of MPI that can run over the network infrastructure provided by the Globus Toolkit
21. *NAMD* - parallel molecular dynamics code for large biomolecular systems
22. *NETCDF* - an interface for array-oriented data access and a library that provides an implementation of the interface
23. *NGS web site* – <http://www.grid-support.ac.uk>
24. *PBSpro* – *Portable Batch System (professional edition)*: a system to organise jobs into different streams and assign different resources and priorities to these streams so as to maintain an optimum mix of jobs on the cluster
25. *R* - a language and environment for statistical computing and graphics. The R environment offers an integrated suite of software facilities for data manipulation, calculation and graphical display.
26. *RUS* – Resource Usage Service: a standardised method of providing usage accounting records using a minimal set of parameters as agreed by the Global Grid Forum (now OGF)
27. *ScaLAPACK* - a library of high-performance linear algebra routines for distributed-memory message-passing systems
28. *SIESTA* - software to perform electronic structure calculations and molecular dynamics simulations of molecules and solids (licence required)
29. *TFBS 0.5.0* - a Bioinformatics framework for transcription factor binding site analysis

