Service Level Definition for UK National Grid Service Resources at Cardiff University

Release 1.0

12 June 2007

1 Introduction and Definitions

1.1 Service

Cardiff University provides two services to the NGS:

1. A cluster of Microsoft Windows XP machines. This is called the **Windows Service**. The Windows Service is non-dedicated and machines are used opportunistically. This means that it is relatively common for a job to be stopped and restarted on another machine. Hence, the Windows Service is most suitable for serial jobs which complete within about 12 hours.
2. A cluster of SGI Irix 6.5 MIPS shared memory machines. This is called the **SGI Service**. The SGI Service is dedicated. This means that once a job starts, it is unlikely to be interrupted. Hence, the SGI Service is suitable for long running jobs, particularly MPI, OpenMP, and multi-threaded jobs.

The term **Service** encompasses both the Windows Service and the SGI Service.

1.2 Service Provider

Cardiff University's Information Services Directorate (INSRV) provides access to the Windows Service from the idle time of laboratory PC's throughout the university. INSRV is the **Service Provider** for the Windows Service.

The Welsh e-Science Centre (WeSC) based at the Cardiff School of Computer Science provides access to the SGI Service. WeSC is the **Service Provider** for the SGI Service.

1.3 Users and Support

The term **NGS User** refers to anyone who obtains access to the Service as a result of the acceptance by the NGS of an application as an individual or as a member of a project. NGS Users must accept the *Regulations for Use of the NGS* available from the NGS website.

The Service Provider reserves the right to restrict or remove access to an NGS User.

The term **User** refers to NGS Users and anyone else who obtains access to the Service through other means. Users who are not NGS Users may be subject to different requirements and limitations to those contained in this document.

A Getting Started document for NGS Users is available on the NGS website.

Specific assistance on the Service should in the first instance be directed to the NGS Support Centre through the Helpdesk contact listed at <http://www.ngs.ac.uk/>. Queries to the NGS Support Centre regarding this Service will be acknowledged within two working days from the time the ticket is assigned. Problem resolution is best-effort: problems will be resolved as and when appropriate Service Provider staff are available.

1.4 Approved Certificate Authority

An **Approved Certificate Authority** is an X.509 or Grid Security Infrastructure (GSI) Certificate Authority (CA) that has been approved by the Service Providers. Services will only accept credentials signed by an Approved Certificate Authority. The UK e-Science Certificate Authority is always an Approved Certificate Authority. Services will normally accept other CA's which the NGS may nominate from time to time. Any exceptions will be notified to the NGS Technical Board and indicated at an appropriate place on the NGS

website.

NGS Users must only access the Service using credentials supplied for their personal use by an Approved Certificate Authority.

1.5 Computation Time

The **Computation Time** is the sum of the actual ("wallclock") time used for each processor that is allocated to a job. Hence, the Computation Time for an individual job is approximately the job runtime multiplied by the number of processors allocated to the job. This means that an 8-processor parallel job will use 8 hours of Computation Time for every actual hour that elapses.

The **Allocated Computation Time** is the minimum Computation Time made available exclusively to be shared amongst all NGS Users over a specified period of time. The amount of Computation Time available for individual NGS Users is not specified. However, the Service Provider will endeavour to ensure that NGS Users with queued computations get a fair share approximately equal to the available allocation divided amongst the total number of active NGS Users.

The Allocated Computation Time may be provided non-uniformly through the time period, and will be forfeited if no suitable jobs are queued to run when the time is available.

Processes that are not started through the batch queue system and processes that fail to complete successfully for any reason will count towards the Allocated Computation Time.

Additional Computation Time may be made available beyond the Allocated Computation Time. The Allocated Computation Time and any additional Computation Time may be provided by alternative resources with a specification equivalent to or better than the nominated resources.

1.6 Software

Cardiff University defines **Core Software** to be software required or requested by all users to provide the basic Service. Support calls relating to Core Software will be dealt with as a matter of highest priority. In-depth knowledge of Core Software is available from staff within the Service Provider. Advice is provided on both basic and advanced use. Problems will be reported to suppliers, their resolution monitored, and fixes applied if supplied.

Cardiff University defines **Supported Software** to be software required or requested by a number of users to provide an enhanced Service. Support calls relating to Supported Software will be dealt with as a matter of medium priority. Some knowledge of Supported Software may be available from staff within the Service Provider whereas some knowledge of Supported Software may be available from staff within other Schools. Advice is provided on basic use only. Problems will be reported to suppliers, their resolution monitored, and fixes applied if supplied.

Cardiff University defines **Unsupported Software** to be software requested by one or more users to provide an enhanced Service. Unsupported Software includes software developed by other suppliers to allow Users to submit jobs with particular requirements such as a job that requires a particular application or database to be installed on the target machine. No knowledge of Unsupported Software is available from staff within the Service Provider.

2 Windows Service

2.1 Hardware

The Windows Service consists of a large number (500+) of university workstations, which are provided by INSRV or other Schools for use in computer laboratories. When the machines are not being used interactively, the spare capacity is made available for batch computation use. Each of the machines in the Windows Service has a minimum specification of:

- 500 MHz Pentium 3 processor
- 256 MB of system memory
- 1 GB of local diskspace
- 10 Mb/s network connection

There is a wide variability in the specifications of machines in the Windows Service. At least 80% of the machines have a specification of at least:

- 2 GHz Pentium 4 processor
- 512 MB of system memory

2.2 Operating System

All execution machines run Microsoft Windows XP (Service Pack 2). Additional patches or service packs or upgrades to new Windows operating systems may be applied at the discretion of the Service Provider. The front-end node for NGS Users of the Windows Service runs Linux.

2.3 Access

Access to the Service is available to NGS Users who have been approved through the NGS application procedure.

Access to the Service for NGS Users is only supported through interfaces secured using the Grid Security Infrastructure (GSI). Supported methods of access are:

- Globus Toolkit 2.4 interfaces to GRAM, GridFTP and MDS. The GRAM interface provides access to a batch queue system as well as an immediate fork capability.
- GSIssh, a version of Secure Shell (ssh) accessible only using GSI credentials.

NGS Users wishing to access the Service using methods other than those listed above must obtain approval from the Service Provider before proceeding.

Accounts may be provided from a pool of generic accounts. If an account is not accessed using one of the above methods for 72 hours, the account may be returned to the pool and all user files deleted.

2.4 Computational Capacity

The Windows Service provides an Allocated Computation Time of 6000 hours per week.

Turnaround time cannot be guaranteed, as the system may be heavily utilised and computations may wait in a queue until a processor becomes available.

The Service Provider may suspend or terminate:

- computations started through the batch queue system that run for longer than 48 hours (measured in wallclock time);
- processes started through any means other than the batch queue system that run

for longer than 48 hours (measured in wallclock time) or longer than 15 minutes (measured in CPU time);

- computations that have the potential to adversely affect the system.

Long-running interactive jobs (i.e. not started by the batch queue system) may have their priority lowered.

Due to the nature of the Condor batch queuing system as a collector of idle machines, the Services are not defined by the number of machines in each cluster. However, as a guide, the Windows Service typically has between 500 and 1000 processors.

NGS Users who require access to capacity in excess of the published limits should obtain approval from the Service Provider before proceeding with their computations.

2.5 Filesystem

The Windows Service does not provide a shared filesystem between nodes. When using the Globus Toolkit to access the Windows Service, an NGS User must define the executable, and any stdin, stdout or data files for the computation as HTTP or GridFTP URLs. Details are provided on the Cardiff section of the NGS website.

Temporary files can be stored on the local filesystem in the “C:\Temp” folder and should be removed when not running computations. The Service Provider may remove NGS User files on the local disk if they are older than 3 days, or if the NGS User has no processes running on the server, or if the files have the potential to adversely affect the system.

2.6 Availability

The Service Provider will attempt to ensure that the system remains available most of the time. However, the system may be unavailable at times due to essential planned maintenance to hardware or software, unplanned outages, or failures out of hours.

If maintenance and/or failures cause the allocated computation time not to be available for any week, additional processor time will be allocated for NGS use in the 4 weeks following the return of the system to fully working order. The total additional processor time available will be limited to one week’s allocated computation time.

If the system remains unavailable for longer than 1 hour, any limits on file lifetime will be extended by the time that the system is unavailable. This does not hold for files in excess of the published limits.

These guarantees do not hold if the lack of availability is due to problems outside Cardiff University (e.g. problems in the JANET network or at the NGS User’s local site).

The reliability of the service is monitored and figures are made available for inspection by NGS Users on the NGS Web site.

Due to the use of pool accounts, there is no guarantee that files will remain on the system, and files may not be backed up.

2.7 Core Software

The Windows Software Core Software includes:

- Condor
- Globus Toolkit

2.8 Supported Software

The Windows Service Supported Software includes:

- Matlab Common Runtime
- Java™ 2 Runtime Environment, Standard Edition (build 1.5.0_06 *or later*)

3 SGI Service

3.1 Hardware

The SGI Service consists of a small number of large shared memory machines. Each of the SGI machines has a minimum specification of:

- 8 x 64-bit 500 MHz MIPS RISC R14000™ processors
- 8 GB of shared system memory
- 12 GB of local diskspace

In addition, the entire SGI system provides:

- access to a 1500 GB SAN Fibre Channel Storage System
- Myrinet™ high-speed message passing interconnects

3.2 Operating System

All machines run SGI Irix 6.5.24. Patches and operating system updates may be applied at the discretion of the Service Provider.

3.3 Access

Access to the Service is available to NGS Users who have been approved through the NGS application procedure.

Access to the Service for NGS Users is only supported through interfaces secured using the Grid Security Infrastructure (GSI). Supported methods of access are:

- Globus Toolkit 2.4 interfaces to GRAM, GridFTP and MDS. The GRAM interface provides access to a batch queue system as well as an immediate fork capability.
- GSIssh, a version of Secure Shell (ssh) accessible only using GSI credentials.

NGS Users wishing to access the Service using methods other than those listed above must obtain approval from the Service Provider before proceeding.

Accounts may be provided from a pool of generic accounts. If an account is not accessed using one of the above methods for 72 hours, the account may be returned to the pool and all user files deleted.

3.4 Computational Capacity

The SGI Service provides an Allocated Computation Time of 1000 hours per week.

Turnaround time cannot be guaranteed, as the system may be heavily utilised and computations may wait in a queue until a processor becomes available.

The Service Provider may suspend or terminate:

- computations started through the batch queue system that run for longer than 48 hours (measured in wallclock time);
- processes started through any means other than the batch queue system that run

for longer than 48 hours (measured in wallclock time) or longer than 15 minutes (measured in CPU time);

- computations that have the potential to adversely affect the system.

Long-running interactive jobs (i.e. not started by the batch queue system) may have their priority lowered.

Due to the nature of the Condor batch queuing system as a collector of idle machines, the Services are not defined by the number of machines in each cluster. However, as a guide, the SGI Service typically has between 24 and 56 processors.

NGS Users who require access to capacity in excess of the published limits should obtain approval from the Service Provider before proceeding with their computations.

3.5 Filesystem

The SGI Service provides a shared filesystem. Each NGS User will have access to a home filesystem on the shared diskspace. NGS Users should limit disk usage on the home filesystem to 40MB while running computations and 10MB otherwise. If an NGS User exceeds these limits, the Service Provider may remove or compress files in order to reduce the NGS User's usage below the limit. If the NGS User does not access the Service for 72 hours using either the Globus Gatekeeper, GridFTP or GSIssh, the Service Provider may remove all files in the account.

Temporary files can be stored on the local filesystem in the "/tmp" directory and should be removed when not running computations. The Service Provider may remove NGS User files on the local disk if they are older than 3 days, or if the NGS User has no processes running on the server, or if the files have the potential to adversely affect the system.

3.6 Availability

The Service Provider will attempt to ensure that the system remains available most of the time. However, the system may be unavailable at times due to essential planned maintenance to hardware or software, unplanned outages, or failures out of hours.

If maintenance and/or failures cause the allocated computation time not to be available for any week, additional processor time will be allocated for NGS use in the 4 weeks following the return of the system to fully working order. The total additional processor time available will be limited to one week's allocated computation time.

If the system remains unavailable for longer than 1 hour, any limits on file lifetime will be extended by the time that the system is unavailable. This does not hold for files in excess of the published limits.

These guarantees do not hold if the lack of availability is due to problems outside Cardiff University (e.g. problems in the JANET network or at the NGS User's local site).

The reliability of the service is monitored and figures are made available for inspection by NGS Users on the NGS Web site.

Due to the use of pool accounts, there is no guarantee that files will remain on the system, and files may not be backed up.

3.7 Core Software

The SGI Service Core Software includes:

- Condor
- Globus Toolkit

3.8 Supported Software

The SGI Service Supported Software includes:

- MIPSpro C compiler
- MIPSpro C++ compiler
- MIPSpro Fortran 77 / 90 / 95 compiler
- Java™ 2 Runtime Environment, Standard Edition (build 1.4.1_06)
- NAG Fortran Library, Mark 21