



THE WHITE ROSE GRID e-Science Centre

Monitoring on the National Grid Service



Introduction

The National Grid Service (NGS) aims to "provide coherent electronic access for UK researchers to all computational and data based resources and facilities required to carry out their research, independent of resource or researcher location". Many of these resources are compute resources, for which monitoring plays a key part in ensuring that they are accessible and available.

The White Rose Grid, through the University of Leeds, supports the monitoring activities undertaken within the NGS in a variety of ways that benefits both the end users as well as the system administrators of the resources made available through the NGS.

Monitoring for Users

The NGS hosts a federated Ganglia service - <http://ganglia.ngs.rl.ac.uk/> - which collects statistics from all of the NGS core sites and certain NGS Partner sites. The Ganglia service provides a wealth of information about the state of the resources, which is useful for system administrators but potentially overwhelming for users.


Ganglia service presents in a way that is much easier to take in. The monitor is situated on the main NGS web page - <http://www.grid-support.ac.uk/> - and updates every minute, live in the user's web browser without the need for manually refreshing the page.

Monitoring for Management

In addition to the tools mentioned above, the National Grid Service also uses Inca (<http://inca.sdsc.edu>), a user-level Grid monitoring framework developed at the San Diego Supercomputer Center. Inca provides a flexible framework that allows the creation of custom tests to ensure that sites are conforming to the requirements set for them. For the NGS, this equates to a suite of tests for the NGS minimum software stack - a set of required tools and interfaces needed to become an NGS Partner or Affiliate site. These include basic tests such as simple Globus authentication checks to full job submission and retrieval via Globus. There are also additional tests for Storage Resource Broker support, compiler version numbers, support for the NGS resource broker and others that sites can choose to have run against their resource.

The tests integrated into Inca are run under a standard NGS user account without system privileges and are run regularly from the four core sites; The Universities of

Load on NGS sites



The NGS load monitor (pictured, left) was produced at the University of Leeds to summarise the load information that the



“The monitoring tools used and created by the NGS allow both users and administrators to get a current picture of the state of the NGS that is relevant to them.”

Leeds, Manchester and Oxford and STFC RAL. This makes it possible to detect problems reliably with few false positives. If a test against a resource fails from all four core NGS sites then it is likely to be a problem with that resource rather than each of the individual NGS core sites. It also makes it possible to detect problems that would not occur for local resource accounts with system privileges. The NGS provides a stand-alone download that allows a site to test itself for conformance to the NGS minimum software stack without needing to add the resource to the NGS testing framework.

The NGS has taken the existing Inca framework and added extra functionality to it beyond the custom tests in order to meet their needs. These enhancements include automatically generated reports detailing pass/failure rates for each resource for the past two fortnights along with graphical representations of results going back several months. The reports are used in the fortnightly NGS Technical board meetings, where the state of NGS provisioning is monitored by the site administrators and managers.

The University of Leeds has been involved in developing the NGS conformance pages available through the NGS Inca site – <http://inca2.ngs.ac.uk/>. This allows access to historical test results, a feature only recently made available in the latest Inca release. This interface shows how a site has performed against the minimum software stack test suite for the past week on a single page. It also provides instant access to error output for any tests that encountered problems. Access to historical Inca results for any one-week period is also available through the same interface. The NGS Inca pages also offer a snapshot

overview of the entire NGS offering a summary of the latest results for each site for every applicable test. The default display mode for Inca lists each individual test result from each site separately, which was considered too verbose for the NGS’s requirements.

Finally, the NGS has developed an opt-in email notification service to ensure that site administrators are informed of problems as soon as possible. Emails are sent directly to the NGS helpdesk to ensure that the information is relayed to the correct people so that action can be taken and any disruption to users can be minimised.

Future Monitoring

As detailed above, the monitoring tools used and created by the NGS allow both users and administrators to get a current picture of the state of the NGS that is relevant to them.

Work will continue in the area of extending and improving additions made to the Inca framework as well as incorporating improvements made available through updates from the San Diego Supercomputer Center. The NGS maintains contact with the Inca development team and participates in Inca workshops wherever possible in order to achieve this goal.

Further Information

Contact:
Dr Shiv Kaushal
(email: s.kaushal@leeds.ac.uk)
National Grid Service:
<http://www.grid-support.ac.uk>
NGS Inca Pages:
<http://inca2.ngs.ac.uk>
SDSC Inca:
<http://inca.sdsc.edu>