

The National Grid Service User Accounting System

Katie Weeks, Kevin Haines, Andrew Richards

STFC, Rutherford Appleton Laboratory, Oxfordshire

Abstract

Monitoring and Accounting are vital requirements for any production grid, used to restrict usage to the allocated quotas and provide statistical overviews of the use of the grid. The UK's Grid, the National Grid Service (NGS), has developed and implemented a User Accounting System (UAS) to automate the processes of user registration and account approval as well as provide automatic policing of accounts that go over quota. The User Accounting System queries the Resource Usage Service (RUS) to obtain usage details for each individual user, inputting the data into an Oracle database which will then send out warning emails to users who reach 90% of their allocated quota and locking accounts of users who reach 100%. Oracle Application Express is used as the web interface to the system, allowing applications to be approved and reviewed, and users to access their account details.

1. Introduction

The National Grid Service is the UK's Grid for academics. There are currently over 500 users from all disciplines including physics, chemistry, molecular biology and the social sciences. Since entering production in October 2004, the National Grid Service has grown from its four core sites at STFC-RAL, the White Rose Grid at the University of Leeds, the University of Manchester and the University of Oxford to include the five partner sites at Bristol University, Cardiff University, Lancaster University, Queens University Belfast and the University of Westminster.

Like any production Grid, accounting is an important part of the monitoring process. Users applying to the NGS declare the number of CPU hours they would like, as well as the amount of storage in the form of disk, SRB or Oracle storage. All applications then undergo a light-weight peer review process to determine the validity of the case for access and the number of hours requested.

In such an environment, it is important to know how much of their allocated quota users have consumed and when they reach their limits. Initially only monitoring was possible, in the form of a CPU check facility that determined the CPU usage of each user, but no accounting process was in place.

As the NGS grows in both the number of users and sites, it becomes increasingly important to account for CPU usage of individuals. Many users will go far over their allocated quotas if not accounted for – even if monitored. It was

required that an automated process could 'lock out' users who went over their quotas, thereby retaining the integrity of the application and peer-review processes.

Accounting is also important for the future, where Grids may introduce a pricing policy for resource usage, whether the price is monetary or not. Pricing policies cannot work without a stable accounting infrastructure in place.

At the beginning of 2006, the National Grid Service had the facilities in place to monitor individual usage as well as the policing policies, but not the tools to use the monitoring data to enforce those policies. It was this problem that the NGS User Accounting System was designed to solve.

2. Other Solutions

Whilst accounting has long been identified as an important feature of a production grid, very few have actually implemented an accounting procedure. An overview of some of the accounting systems can be found in table 1 [1].

Name	Developed by	Used by
DGAS	EGEE	EGEE
APEL	LCG	LCG/EGEE
GSAX	GGF	None
RUS	GGF	NGS
SGAS	SweGrid	SweGrid, Globus4
GASA	Gridbus	Gridbus
AMIE + TGAcounting	TeraGrid	TeraGrid

Table 1: Grid Accounting Systems Overview

2.1 EGEE

APEL

The LHC Computing Grid (LCG) developed the implementation APEL (Accounting Processor for Event Logs) using the R-GMA (Relational Grid Monitoring Architecture) [2]. Accounting information such as CPU time, Wall Clock time and user Distinguished Name is collected via the Distributed Grid Accounting System (DGAS) and published via R-GMA. The APEL application then processes the log information to create true accounting records, publishing the information back to the R-GMA where it can be collated at the Grid Operations Centre (GOC). Because of the Virtual Organisation (VO) emphasis within EGEE there is no automatic notification or account deactivation when a user goes over quota. Quotas are managed on a VO rather than individual basis.

2.2 TeraGrid

TeraGrid in the US allows users to monitor their usage via a command line interface, but have no architecture in place to allow for automated accounting, account deactivation or automatic notifications [3].

2.3 Open Grid Forum

Resource Usage Service

Since 2003 the Open Grid Forum has been developing the Resource Usage Service (RUS) [4, 5]. The RUS stores monitoring information by linking jobs to users. This information can then be queried by other services wishing to use monitoring information to perform accounting on individual users.

2.4 The NGS Solution

The National Grid Service has had a Resource Usage Service in production since early 2006. It was decided to continue using the RUS and design the NGS User Accounting System to interface with the NGS RUS. The RUS already provided every feature the NGS needed and not continuing to use it but changing to another accounting and monitoring system would not have added any extra functionality.

3. The NGS User Accounting System

Figure 1 gives an overview of the components and interfaces of the NGS User Accounting System. The accounting Oracle database lies at the core of the User Accounting System. Information such as the user's details, the application information, the account status, the

resource allocation and the resource usage is stored there. For the full database schema see Figure 5 at the end of the paper. The resource usage data is obtained by querying the Resource Usage Service (RUS) [4, 5] at Manchester once a day and uses the collated usage over all NGS sites to populate the database.

The database also maintains historical information about resource usage, so that usage can be tracked over a period of time. It also maintains information about any changes to a user's certificate Distinguished Name (DN). This is important as the system needs to account for usage by user, not simply by DN. The information in the RUS is based on the DN of the user who submitted the job, and this DN may change over time, for example as the user moves institutions.

The design aims of the NGS User Accounting System were to automate the application process, provide the tools for automatic policing of accounts and to provide a system that could incorporate new functionality as the NGS develops. The systems in the diagram (website, helpdesk, RUS database, LDAP/VO database and SRB/JISCmail management) were all in place prior to the development of the UAS, which gathered all the systems together.

3.1 Application process

When a user applies to the NGS using the application form on the NGS website (1), the details they give are stored in the database under their Distinguished Name (2). An account is created for the user with a pending status, and no resources allocated. An email is then sent to the NGS helpdesk to inform the system administrators that a user has applied (3), and the peer-review process starts immediately after a member of NGS staff has changed the helpdesk ticket status to 'under review'. Once the application has been peer reviewed (4) and approved (5), the system administrators access the User Accounting System (UAS) via a form that allows them with one-click to approve a user; allocating their resources, updating the status of their account to active, creating their SRB account (7), adding them to the JISC mailing lists (7) and sending an email to the user informing them their request was successful (6). The UAS, when approving a user, sends the DN to an external system (7) that adds the DN to the sites' gridmap files, creates the SRB account and sends out emails requesting the users addition to the JISC mailing lists.

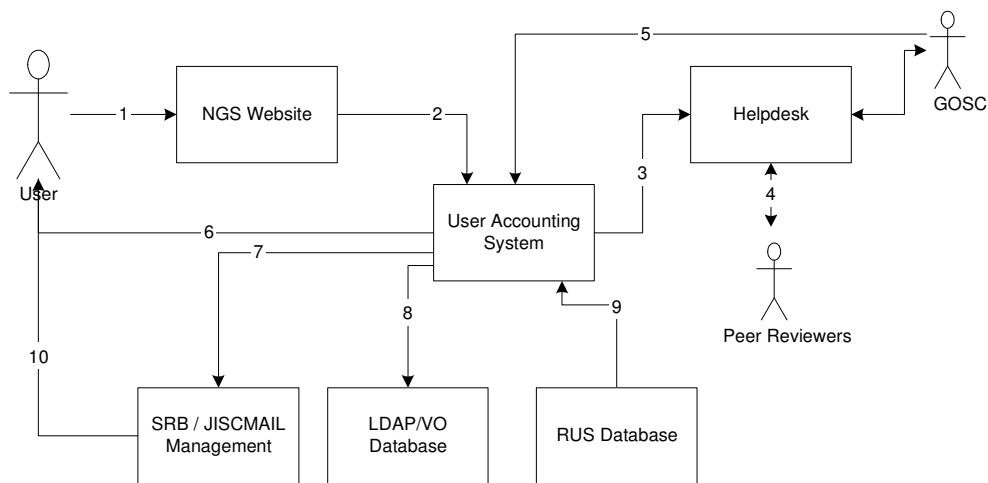


Figure 1: The NGS User Accounting System Interfaces

3.2 Policing

To manage the access of users to the core sites, the LDAP (Lightweight Directory Access Protocol) server queries the Accounting Database every hour (8). The LDAP is then populated with only the Distinguished Names of users who have an active account. This means that as soon as an account is unlocked, it will be only a maximum of an hour before the account is active on all nodes again.

Once a user has an active account on the NGS, they come under the accounting policies of the NGS User Accounting System. The UAS queries the RUS (9) every day for total CPU hours and disk space for each user. This information is then checked against the allocated resources, and the UAS sends a warning email to the user if they are over 90% of their quota (6). The users account status is updated to show that the account is under warning.

If a user's account goes over 100% of quota for CPU hours, the UAS locks the account, the users DN is removed from the LDAP server (8) and the user is emailed (6). The user is then invited to apply for further resources using the original application form or via their account details page. If the renewal application is successful, the user's resource allocation is updated. This change to the user's allocated

resources in the Oracle database causes a trigger to fire that re-evaluates the resources used against the new allocation, and sends an appropriate email to the user (6). The account is then unlocked, the warning flag removed from it and an email sent to the user to inform them they are within limits and active again (6).

4. Integration with the NGS Helpdesk

The NGS uses Footprints as its helpdesk system. The application review process was originally dealt with via Footprints and as such it was decided, for ease of user for reviewers, to keep it within Footprints. This decision required that the system send the application directly to Footprints, from where the review system could be unchanged from the reviewer's point of view. However this has the unfortunate consequence of dividing the approval process into two parts, with all applications requiring approval in the User Accounting System database as well as in Footprints so the ticket can be closed.

This synchronisation of the two systems is currently done manually and is a far from ideal situation. Future considerations for the NGS User Accounting System include removing the Footprints dependency and integrating the peer-review process into the System.

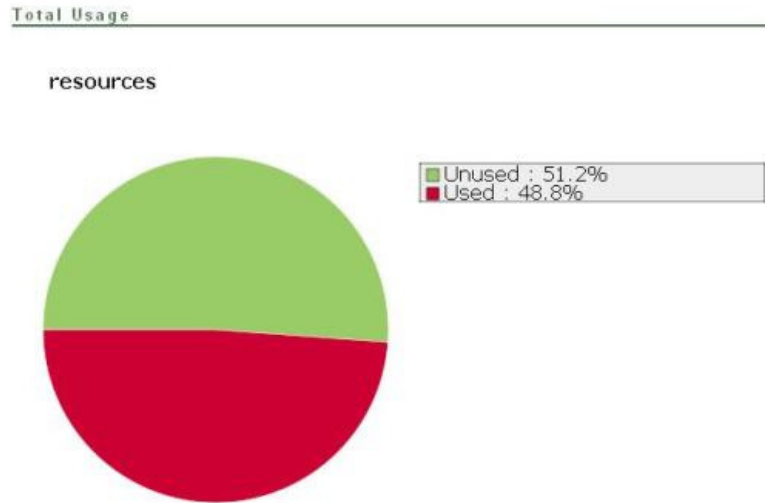


Figure 2: CPU usage as percentage of allocated limit – detail of what users can see for their account

5. Installation and Commissioning

When the NGS User Accounting System was initially bought into production, there were a number of users who were not only 100% over quota, but up to 10,000% over quota. To prevent a large number of users being locked out in the initial stages and to assist in the task of debugging the system, the top users only were locked out in groups of ten at a time. As the system was unfamiliar to users, this also helped the helpdesk deal with the queries that locking accounts generated.

The method of only notifying the top ten users at a time was successful, with only a few queries to the helpdesk and 60% of users applying for renewal of their resources.

The NGS User Accounting System is now policing at production levels with account deactivation at 100% and warning emails being sent when a user reaches 90% of their allocation.

5.1 Administration Functions

The UAL has facilities for user account management by the NGS administration staff. Administrators can lock and unlock user accounts (the system will not override a manually locked/unlocked account), edit users' resource limits, and view their application and renewal justifications. They can also alter all account details such as email addresses and Distinguished Names.

5.2 User Access to User Accounting System

As users got used to the new system and being accounted for, they began to request easy access to their accounting records and account details. Using browser-based X.509 certificates, users are now able to access their own account details. This makes monitoring their accounts much easier for individuals, as they are able to see at a glance in report and graph form their quota and usage (Figure 2). Other features include being able to change contact details, checking the status of their account and when they applied, and renewing resources from there rather than the main application form.

5.3 Problems Encountered

The use of Oracle Application Express posed a problem in a few areas. Whilst it made the design of forms and reports simple and quick, it also imposed limitations. X.509 certificate access is not supported by Oracle Application Express and so a work around in the form of a PHP script was implemented in order to allow users to access their account details by certificate authentication, using the most user-friendly method of retrieving certificate details from the user's browser.

The length of some of the fields on the application form also caused difficulties at first, with application case studies being limited to 2200 characters, a problem when more detailed information was required to approve some of the renewals.




Figure 3: The number of new applications to the NGS since entering production

6. Results

Since entering production in October 2006, 57 users have had their accounts deactivated due to exceeding their CPU quota. 66% of those who were locked out applied successfully for further resource allocations. No reasons are given for those who did not renew their accounts, but they could be situations such as PhD students who used the NGS during their studies but who have now left and not continued in research.

Since October 2006, 80 new users have applied for NGS accounts (Figure 3), with the User Accounting System handling the automatic registration process successfully for all of them.

Real-time statistics have been developed as part of the User Accounting System. The use of Oracle and Apex reduced significantly the time taken to produce statistics such as CPU usage across core sites (Figure 4), and also allowed the production of statistics such as where users heard about the NGS.

The development of the NGS User Accounting System is felt to have exceeded the design aims and its ability to grow and develop with the NGS will ensure it is a robust, long-lasting system.

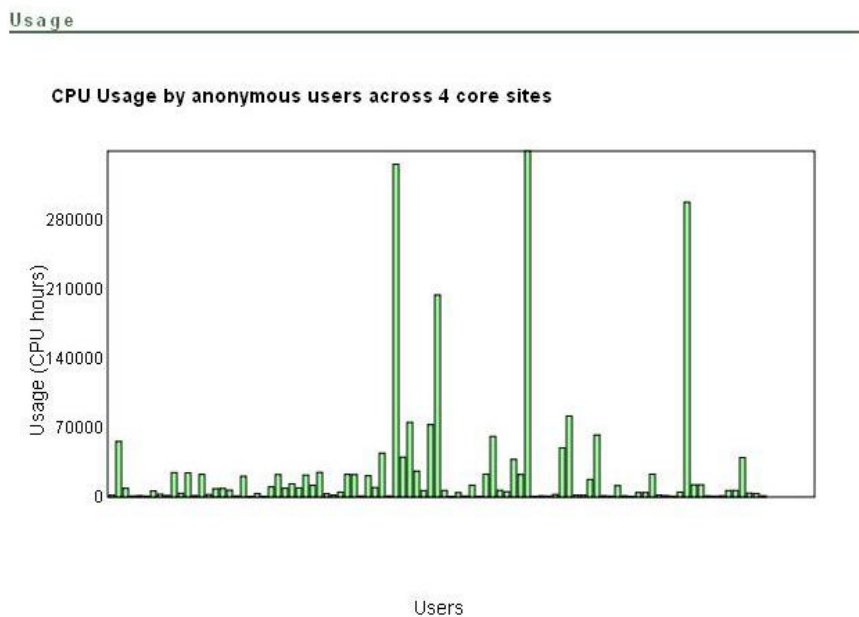


Figure 4: Example of the statistical graphs produced by the NGS User Accounting System

7. Future Work

Work on the NGS User Accounting System is ongoing and there are a number of improvements next on the agenda. The peer-review process, currently being handled in Footprints, is under consideration to be moved into the Oracle system. The system also currently links to an external system that auto-creates an SRB account when each user is approved, as well as sends registration emails to the JISCmail NGS mailing lists. This residual connection to an external system needs to be taken into consideration when discussing the future developments and requirements of the NGS User Accounting System.

Another area for development is in collecting data on disk/storage usage by the users, which is presently not done. The system is designed to be easily expandable to monitor any type of resource. The NGS User Accounting System currently polices only CPU time. When both Disc and SRB storage are able to be accounted for, the UAS can be simply extended to include policing of resources other than CPU hours.

The simple extension of the NGS User Accounting System to support Virtual Organisations is expected as the NGS adopts VO technology.

8. Conclusion

The development of the User Accounting System for the NGS has made accounting a much simpler procedure. Users are now automatically notified when they are at 90% of their allocated CPU quota, and accounts are locked when they go over 100%. No manual intervention is necessary - once locked out users can apply for more time and their request goes through the standard application process.

The application process has also been simplified by the introduction of the User Accounting System, with user details being entered into the database upon application and the approval process becoming more automatic.

Access to individual user accounts is made simple for all users by the browser-based certificate access to their accounts, where they can alter contact details, see the status of their accounts, apply for further resources and check their usage.

The introduction of the User Accounting System has also made it possible to create real-time statistics such as number of applications per month and where users heard about the NGS. Further additions to the User Accounting System will see more statistics being produced.

