

Grid-SAFE
Resource Usage Monitoring and Accounting
Stephen Booth, Adrian Jackson, Nicola McDonnell, Josh Green, Davy Virdee
EPCC, University of Edinburgh

This extended abstract outlines a paper which discusses the work that has been undertaken in the Grid-SAFE project (a JISC funded project) to provide a set of tools/services for resource usage monitoring and accounting enabling users and managers of HPC and Grid resources to perform these tasks in a standard way.

Resource usage monitoring and accounting has long been standard practice on high-end compute resources, such as the national HPC facilities. Such compute resources are expensive to procure and run, and as such necessitate the need to be able to re-coup or assign costs to individual users, groups, projects, institutions, or other such groupings.

Such usage monitoring and accounting requires three distinct sets of functionality to provide good quality service. Users - that is managers of the resources who want to undertake the accounting and monitoring operations, or resource users who want to monitor how much resources they've used - need to be able to collect the usage information from the service. Once the usage data has been collected there needs to be a mechanism for processing and adding 'policy' based information. This process converts the raw data from the resource into something that references usable concepts and can be mapped to the specific requirements of the service/facility being provide; for instance it may be mapping usernames on a HPC machine to people that can be identified in the real world, or recording that a particular piece of usage data is associated with a particular piece of hardware or charging model. Finally, once the processing of the data has occurred the user requires some mechanism for querying the data and producing usable and useful reports.

Whilst it has long been a requirement to undertake this type of activity for high-end computing resources, there is a lack of tools or products available that provide all three functions. The job submission and batch control software that generally manage users' jobs on these systems provide large amount of information on resource usage, and it is generally a straightforward task to collect that information and store it in an accessible form (often a database is used). However, the processing and reporting functionality is often performed either by hand, or with bespoke pieces of software. Such software may fulfil the requirements for a particular resource or user, but cannot be reused for other resources or for other types of reporting or policies.

Furthermore, when considering a more complicated environment than the standard high-end computing resource (where there is generally one machine with a distinct set of policies) such as computational Grid; the task of collating, processing, and reporting on resource usage information becomes more difficult. There may be a number of resources, often of different types and using different software, at different locations. It is usual for there to be a number of distinct user groups (i.e. virtual organisations) and management groups (i.e. Grid managers, local resource managers and managers of each user group). This means that the sources of data, types of resources, and required policies for the Grid scenario can be much greater than the traditional high-end computing resource.

EPCC, through the Grid-SAFE project funded by JISC, have been addressing this lack of usage monitoring and accounting tools. Building on the success of our SAFE (Service Administration

From EPCC) tool which was developed to fill this tool gap for the UK national HPC services, HPCx and HECToR, we have been working to create a customisable monitoring and accounting tool that can be used by local HPC managers, campus computing managers, national HPC resources, and Grids alike.

The Grid-SAFE project follows two main strands of work; Accounting Framework and Client/Usage Monitoring, which have been progressed since May 2008. Developing the accounting framework has involved actively working with the OGF¹ RUS (Resource Usage Service) working group to progress the standards to provide usable and useful service definitions and standards for accounting and resource usage services. This has culminated in the production the Resource Usage Publishing Interface (RUPI) specification: an interface designed to allow usage data to be uploaded to a service. This interface forms a subset of the functionality provided by RUS and is based on data in the OGF UR (Usage Record) format. Such an interface is crucial for allowing interoperability between the accounting and reporting services being developed for Grid-SAFE and other data repositories, machines/resource or management tools that are already in operation. All that is required is an implementation of a UR provider service to take existing or new data and map it into UR format which can then be processed by a RUPI service. Grid-SAFE provides both UR providers (i.e. a tool that allows providers to generate OGF-UR records for their resources) and a RUPI service.

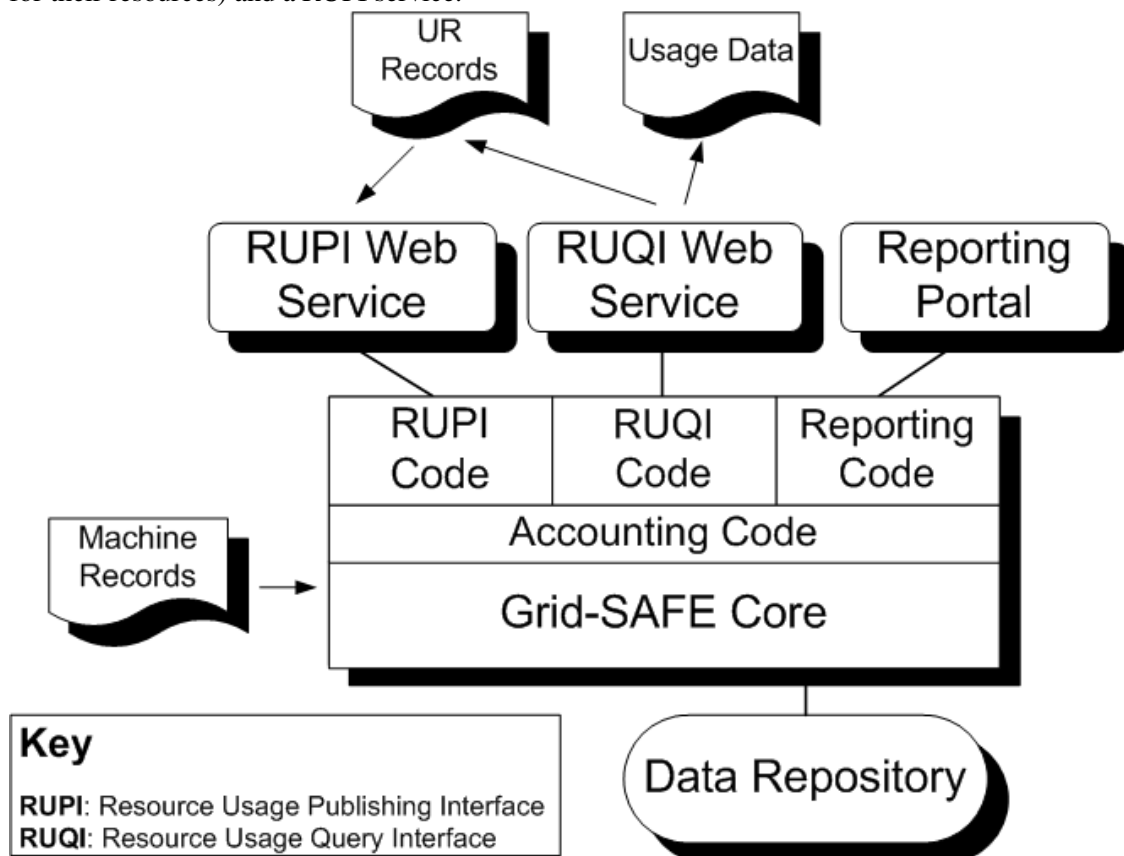


Figure 1: Grid-SAFE general architecture

¹ OGF: Open Grid Forum: <http://www.ogf.org>

The Grid-SAFE project has also been developing an opposite interface to RUPI; the Resource Usage Query Interface (RUQI). This interface is designed to allow users or other services to query a service (such as Grid-SAFE) for usage data.

Furthermore, we have developed a reporting portal that enables users/administrators to generate reports on usage/accounting data. This may be done by using some pre-defined report templates or via the production of bespoke reports to query the specific data required.

Figure 1 outlines the general Grid-SAFE architecture and shows how the services interact. Grid-SAFE is designed to manage HPC machines locally, as well as providing interfaces for data from remote resources. Therefore it allows machine records (or usage data) to be added directly, as well as through the RUPI interface, and it allows querying by remote services as well as letting users generate reports locally (i.e. via a web interface).

Whilst the architecture of Grid-SAFE is straightforward the way it is used and deployed can be very varied. The traditional deployment scenario has been to set it up at the same location as the HPC resource(s) being monitored. Reporting is then performed through a web based system. However, with the new interfaces and tools that have been developed, Grid-SAFE can be deployed with no associated HPC resources and just be used to store and query usage records generated elsewhere, or can be used to report on usage data stored and maintained elsewhere. It is also able to interface with Grid authentication mechanisms such as VOMS and thus provide accounting and reporting for a range of different Grid environments.