THE WHITE ROSE GRID
e-Science Centre of Excellence

Community Resource Brokering on the White Rose Grid

The White Rose Grid provides an infrastructure for experimenting with virtual resources of many types (compute, storage, software, networking etc.) as Grid services. Therefore the consolidation of resource management is required as the resources are spread across geographically distributed domains.

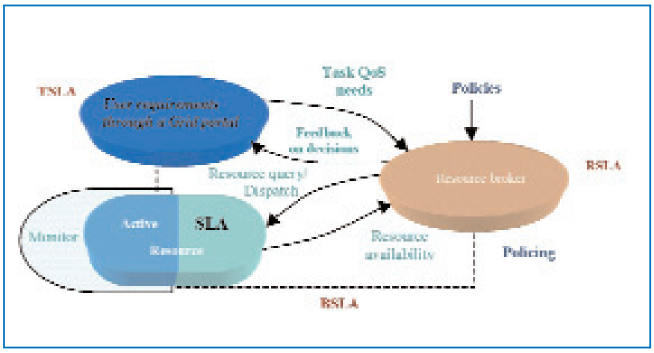

Figure 1: Overview of the SNAP based Resource Broker

To support application execution in the context of the Grid, scheduling Grid resources is necessary. Grid Resource Scheduling (brokering) is defined as the process of making scheduling decisions involving resources over multiple administrative domains. This can include searching multiple administrative domains to use a single machine or scheduling a single job to use multiple resources at a single site or multiple sites. A Grid resource broker must make resource selection decisions in an environment where it has no control over the local resources, the resources are distributed, and information about the resources is often limited or dated.

The community resource broker on the WRG insulates the user from Grid middleware, enabling transparent submission of jobs to the Grid (Figure 1). The broker takes the user's requirements (e.g. number of CPUs, operating system) and job description and contacts resources that may support these requirements to gather information on their current state (e.g. current load). A decision is made as to which resource(s) will be used to run the job and this is followed by a negotiation with these resources. This negotiation is based on the framework provided by the Service Negotiation and Acquisition Protocol (SNAP), whereby guarantees are obtained that the user's requirements will be fulfilled by the use of a Service Level Agreement (SLA). SNAP comprises three main SLAs namely Task Service Level Agreement (TSLA), Resource Service Level Agreement (RSLA) and Binding Service Level Agreement (BSLA).

TSLA is where the user provides a clear objective specification of his task requirements and any resource preferences.

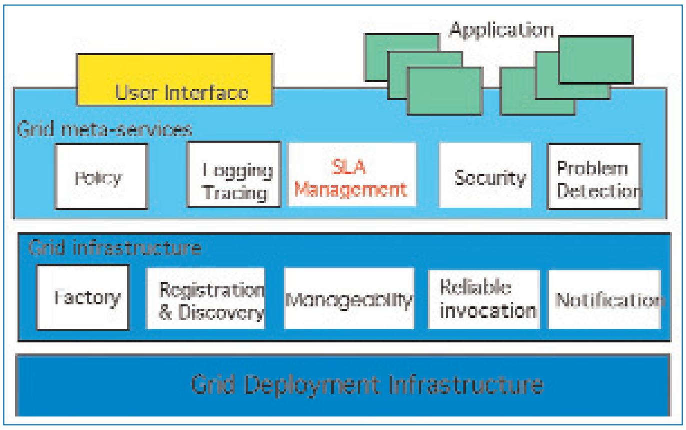

Figure 2: SLA Management

The requirements are ascertained through an interface such as the Grid portal that forwards the users description to the broker for processing. RSLA relates to resource discovery and decision-making on the appropriate resources that meet the task's requirements and securing the resources for utilisation. BSLA only associates the task with the resources.

Users are also bound by commitments as specified in terms of SLAs. An SLA determines the "contract" between the user and the Grid Service provider. It provides the basis for a Grid service in which users can use Grid application service providers to access services and with quality of service agreements enabling them to meet their business commitments. At any given point in time hundreds of SLAs may exist, and each SLA in turn may have a large number of metrics to be observed. A formalised representation of commitments in the form of SLA documents is required so that information collection and SLA evaluation can be automated (Figure 2). SLAs are distributed, and their validation depends on local measurements. For example, for a compute service the resource broker takes the user's requirements and job description then contacts resources that may support these requirements to gather information on their current state.

Accountability and monitoring become increasingly important. Thus appropriate tools are developed for observing resource utilisation, managing performance degradations, availability and other parameters. Observed data need to be validated against commitments made to users. Automation of these processes is important due to the scale and complexity involved.

For these services to be provided reliably and on demand, an organisation such as the WRG needs to have payment mechanisms for providing these resources. By integrating the ability to charge for Grid services, new models of service provisioning such as utility computing are enabled. The development of an infrastructure to enable the trading of Grid services on the WRG is critical and fits within the Grid Economic Services Architecture (GESA)[1,2] being developed within the Global Grid Forum (GGF)[3]. Charge models will be put in place to ensure a consistent and supportable costing practice for the WRG users.

References

1. Grid Economic Services Architecture (GESA-WG) Working Group: http://www.ggf.org/3_SRM/gesa.htm
2. Grid Economic Services Specification: http://www.doc.ic.ac.uk/~sjn5/GGF/draft-ggf-gesa-services-1.pdf
3. Global Grid Forum: http://www/ggf.org

Further Information

Contact:
Dr K Djemame,
School of Computing, University of Leeds,
email: karim@comp.leeds.ac.uk