

A Market for Computational Services

A Proposal to the e-Science Core Technology Programme

by

**The London e-Science Centre
The CLRC e-Science Centre
The North West e-Science Centre
The Southampton e-Science Centre
The University of Wales at Swansea
The Astrophysics Research Institute at
Liverpool John Moores University**

in collaboration with

**Computer Sciences Corporation
Numerical Algorithms Group
GridX
RealTime Engineering Ltd.
Sun Microsystems Ltd.
SGI Inc.
Oracle**

1. Introduction

The intention of Grid computing is to increase the availability and utilisation of computational resources or services. In computational services we include all those elements that can be put together to perform a useful activity: scientific instruments, processors, storage, network bandwidth, software and data. Current developments in Grid middleware: Globus and the Open Grid Services Architecture (OGSA) are providing mechanisms for service registration, discovery and use. However there is currently no widely accepted mechanism whereby charges can be made for the use of these services and payments received. Such mechanisms are essential if the Grid is going to develop from its current experimental status into professional or commercial use. If effective lightweight mechanisms for accounting, charging and payment can be developed it is likely that an open market in Grid services would develop, co-ordinating users and providers of services. Given that such a market would be global and electronically mediated it is likely that many of the theoretical benefits of the free market would actually be realised, greatly encouraging the development and use of Grid services.

We therefore propose a project to develop contracting, accounting, charging and payment mechanisms for Grid services and to investigate the brokering and market mechanisms that could be developed once these facilities are in place. This project is aimed at producing professional accounting, charging and payment mechanisms and software for the UK e-Science Grid and to provide UK input into the OGSA, its open source reference implementation, the Globus Toolkit 3.0 (GT3), and the Global Grid Forum (GGF) through the use of emerging Web Services technologies.

2. Background

The realisation of a full open commercial market in Grid services requires the development of the following capabilities:

1. Mechanisms for service definition, development, registration, discovery and use under defined policy and security constraints.
2. Agreed protocols which services can use to advertise their capabilities and charging policies to Grid users or software agents in such a way that costs of resource usage can be used for negotiating and agreeing a service contract covering expected performance and cost
3. Mechanisms for recording service use and accounting this use to individuals, groups or organisations.
4. Secure mechanisms for charging and payment.
5. Brokering mechanisms to optimise service selection according to a user's or application's need, incorporating both performance and cost information as defined by the protocols in (3). This area, Performance Engineering, is one where the UK has particular expertise and experience.
6. Full market mechanisms to mediate between users and providers of computational services. For example trading, auctions or financial instruments (futures). Stochastic optimisation techniques (games) to maximise collective utility for both users and providers of services.

The first is the active concern of the current Grid middleware activities (Globus and OGSA through its leverage of the Web Service infrastructure). The other topics are not currently receiving as much attention or being closely integrated with the development and standardisation of Grid services. Charging and payment mechanisms must apply not only to hardware resources but also, for example, to software and data. Moreover a transaction should be possible between a user and provider of a service without the need for any prior contract being established as long as access policy conditions are satisfied.

It is our belief, therefore, that charging and payment mechanisms must be lightweight and integrated with the service access and policy mechanisms. We anticipate that a user may wish to compose several

independently sourced services to accomplish a particular task, for example composing a software component from provider A with another software service from provider B, using data owned by C, executing the composed application on a machine owned by facility provider D and utilising guaranteed QoS network capacity from provider E. We therefore, ultimately, envisage a “pay-by-use” model, where appropriate payments would be automatically remitted to providers A, B, C, D and E according to the use made of their services on each invocation. This view seems compatible with Web Services and would facilitate the development of the next generation commercial Internet or Web, where not only information but also useful computational services could be available globally.

It is unlikely that a single charging model or a single basis for charging will emerge in a diverse community. It is therefore necessary to develop a mechanism whereby different policies and clients can meet, i.e. to create a *market*. We agree with a remark by Miron Livney that the only way to realise a true market is ultimately to base it on real money. Once mechanisms for charging and payment for computational services are in place it is interesting to speculate on the effects this would have, particularly on national HPC provision. Thus Research Council funded research groups could be afforded access to HPC facilities by being given funds to purchase HPC services, just as they are currently given funds to hire staff or to travel. These services could then be purchased from central HPC suppliers in much the same way as now but if independent suppliers could provide equivalent services more cheaply they could be used. It should be noted that this is critically dependent on the existence of the protocols in point (2) above. There is no agreed definition of what “computational power” could mean analogous to the agreed definition of “electrical power” used in a utilities grid. Since computational resource consumption is much more complicated than consumption of energy this question needs to be carefully considered.

The market dimension is important as, coupled with the adoption of Web Services, it provides a route whereby the Grid paradigm can be extended to uniformly encompass scientific, commercial and social uses of computation. It provides an arena where not only established practitioners in HPC can develop new, more flexible, ways of delivering their services but also provides opportunities for new small-scale commercial entrepreneurial operators to enter the HPC market. This consortium therefore includes representatives of all such players, not only established HPC providers such as CSC, Sun, SGI, Oracle and Nag but also smaller SME’s and start ups such as RealTime Engineering and GridX.

3. Programme and Deliverables

3.1. Project Overview and Goals

The creation of a market requires the development of a framework, standards and mechanisms that will accommodate all potential players. Thus the focus of the project will be to produce solutions that are not only feasible technically but which will also support realistic business models for all vendors, are compatible with current and planned software infrastructures and will be adopted as international standards. The consortium includes major representatives of all stakeholders in this enterprise: service vendors, infrastructure developers and Grid researchers. A major thrust of this project, therefore, will be to take advantage of this forum to ensure that any proposals developed are both practical and adequate. Not only compatible with existing and planned infrastructures but also capable of supporting a commercial market in Grid services and enabling ambitious long-term research in the development of Grid economics. The process is very much one of defining the “neck of the hour glass”. That is the consortium will draw on its wide experience to identify the services and protocols necessary to support the complete market vision, these protocols will then be specified and a prototype implementation produced – the neck of the hour glass. This will provide the basis on which more general work can be conducted both to test the practicality of these solutions via realistic trials and to develop more sophisticated brokering and trading mechanisms.

The project therefore breaks down into four major stages:

Stage I. Identification of Service Framework and Protocols

All the partners will work together to define future business models through typical use cases and to specify protocols for the basic services of negotiating, recording, accounting, charging and payment

covering all modalities of service (WPs 1 & 2). Having defined the functionality and behaviour of the core services and protocols within a computational economic framework these will be documented to serve as input into a standardisation process through GGF and other relevant bodies (WP 9).

Stage II Implementation of Core Services

Following the requirements analysis and specification conducted in Stage I we will commence implementation of the core OGSA compatible services. The Grid Accounting service will be developed for a beta release with GT3 (WP3). The other services, Grid Banking and Grid Brokering, will be prototyped for release as additional complimentary services deployable within GT3 (WP 4 & 5).

Stage III Testbed Trials

Trial deployments of the core services will be made in three areas (WP 7):

- ❖ Computer Sciences Corporation (CSC) will carry out a series of trial deployments. Their first deployment will be as an alpha test within one or two of their client HPC installations to check the practicality of the resource accounting, brokering, security and system management functionality. Subsequently they would connect two or more such Grids, including their e-HPC facilities, to test full accounting and invoicing.
- ❖ UK e-Science Grid. Here the core functionality will be added to the UK Grid middleware deployed at selected Regional Centres, including those participating in this project, and experiments in resource sharing and accounting conducted, motivated by an exemplar e-science application (WP 6).
- ❖ The Astrophysics Research Institute at Liverpool John Moores University will implement the market services within the eStar e-Science demonstrator and conduct an eStar trial deployment to overseas telescope sites run by UK organisations. These will be arranged in Spain and the US, providing a test of a transnational scientific Grid economy.

Stage IV Brokering and Performance Engineering

Once the kernel services are in place and verified there will be considerable opportunities to develop more sophisticated brokering and trading mechanisms that will 'plug in' to the established framework. Within the project we propose to develop a prototype implementation of a brokering infrastructure that is capable of taking performance, cost and policy information into account to negotiate, acquire and pay for resources.

Richer market mechanisms, such as the creation of instruments in computational resources and the exploration of the consequences of the creation of an open computational market, will be explored in other projects but the requirements to support these advanced mechanisms will strongly influence the standards decided in Stage I. Similarly the industrial partners will be conducting work outside the project to produce middleware and products compatible with the standards agreed within the project.

The consortium, particularly the industrial partners, will work closely with organisations such as the UK e-Science Directorate and the GGF to ensure that the protocols defined in Stage I and demonstrated in Stages III and IV are adopted as international standards and become part of the OGSA.

This work will be carried out in tandem with other work within GGF towards computational economies. It will provide a basis for a proposal for the resource costing protocols to be submitted to the GGF standards adoption procedure. It will be an important part of Stages II and III to provide the experimental validation of the viability of any protocol proposals.

3.2. Work Packages and Deliverables

Work Package 1: Identification of Economic Models and Use Cases

Discussion will be held amongst all partners to identify a feasible economic model that covers all modalities of service and is compatible with the OGSA and commercial Web services. This process will be informed by input from all the partners on their business needs or plans, infrastructure development and research requirements. This work will benefit from the experience and expertise of the following organisations:

CSC, NAG, Sun, SGI and Oracle in providing computational HPC, software and data services,

the eStar team in providing a scientific instrument service,

RealTime Engineering in providing service characterisation, contract negotiation and payment expertise,

Sun and SGI in providing Grid infrastructure and vendor experience,

and the e-Science Centres providing the long-term research requirement.

The e-Science perspective will be provided by all the academic partners but specifically by the requirements of the GEODISE project and Manchester's recent investigation of brokering techniques part of the EU GRIP project.

This work will be captured in several use cases that will reflect the potentially different needs of the commercial and academic collaborators that can be met by a common service infrastructure. The results of this work will be critical in specifying the overall architecture.

Deliverables:

The use cases will be disseminated through a public web site and contributed to the relevant working groups within the GGF (see WP 9).

D1.1: The Application Service Provider use case.

D1.2: The Computational Service Provide use case.

D1.3: The Scientific User use case.

D1.4: The Resource Reseller use case.

D1.5: The Component Supplier use case (that is, a supplier of components that are embedded and used, as required, by one or more of the above but typically without any direct involvement between the Component Supplier and the end user)

D1.6: The Instrument Service Provider use case

D1.7: The Utility Computing use case.

D1.8: The Data Provider use case.

Work Package 2: Protocol and Service Specification

The use cases generated from WP1 will motivate the development of the protocols and service specification for the economic interaction of hardware, software and data services with a consuming client. These will all need to support (in some form) the following processes:

- Advertising. Making public the characteristics of a service and its capabilities and cost. For many types of service, e.g. processing capacity, there is not a uniform commodity and resource users or brokers must be given the ability to compare similar services that are described differently.
- Price and contract negotiation. A service must be capable of interacting with a user or a broker to negotiate a price and contact for the performance of a service and provide mechanisms for

monitoring the performance of this contract and, if necessary, reacting in real-time to changing circumstances which may affect the delivery of the contract.

- Recording and accounting. A service must be capable of recording its usage by a particular user in the appropriate dimensions, storing this information and deriving appropriate reporting for individuals, groups and organisations.
- Charging and receiving payment. A service must support mechanisms for automatically charging for its use and ensuring payment is remitted appropriately.

An architecture capable of supporting all these processes will be developed and the necessary interfaces and protocols defined. This could include the identification of an extensible service definition language, or if no such thing exists, the construction of one. This language must be able to express resources and resource requirements in various degrees of abstraction,

The service interfaces and protocols identified in this work package will be documented and disseminated to the wider community for comment and standardisation through the relevant working groups within the GGF (see WP 9). Reference documentation for these protocols and service interfaces will be maintained on the project web site as they are developed and enhanced over the course of the project and through discussion within the GGF.

Deliverables

We will deliver detailed specification of the service interfaces and associated specifications to support:

D2.1: The use of the OGSA serviceData model to define an ontology relating to **resource capability** encapsulated within an XML schema. Such schemas will be critical in comparing the capability of different services. The temporal validity of this information within the serviceData model can be encapsulated within the existing OGSA infrastructure.

D2.2: The use of the OGSA serviceData model to define an ontology relating to **negotiation mechanisms** encapsulated within an XML schema. This information will be used to derive the negotiation mechanism used by the Grid Brokering service to set the cost for the resource use. The temporal validity of this mechanism can be encapsulated within the existing OGSA infrastructure.

D2.3: A Grid Accounting Service will be defined to enable the recording of resource utilisation by other grid services. This service will also provide mechanisms for retrieving accounting information by properly authorised entities expressed within a local access policy.

D2.4: A Grid Banking Service will be defined to enable secure financial transactions between grid entities. Each entity will have a unique identity within the bank to record all financial transactions.

D2.5: A Grid Brokering Service will negotiate service contracts on behalf of a resource consumer and resource provider for resource utilisation. It is expected that brokering will have both pre-runtime and runtime aspects: pre-runtime to determine the service contract with the User and runtime to determine the optimal way to deliver the contract given the potentially dynamically changing computational and software resources available.

D2.6: A Chargeable Grid Service interface will be developed to encapsulate existing Grid Services for use within a computational economic framework. This interface will allow the resource provider to specify and enforce the services usage and charging policy. Its basic functionality will capture the actual resource usage and on completion of the service invocation initiate the financial transaction between the consumer and provider. This interface will be developed in close collaboration with the GGF OGSI working group, as it will ultimately be migrated into the core Grid Service Specification.

Strawman Architecture

To clarify the control and data flow between these services we present the basic system architecture. The Grid User contacts a brokering service (1) specifying their functional service requirements and any associated constraints. The brokering service extracts economic service meta-data from the OGSA

Computational service (2), which is supplemented by functional data from the underlying Grid Service (3). Having selected a service (4), the user's use is verified against an established contract before recording service initiation and invoking the underlying service interface (5). On completion the resource usage is recorded (6) and the actual cost of the service invocation (established from the accounting records and contract information) is passed onto the banking service (7). The result from the service invocation is passed back to the client (8).

As control flows between each stage of this architectural model logging events are capable of being generated. Other services within the OGSA (such as Auditing Services) that wish to be notified of these events may do so by subscribing through a standard OGSA notification interface.

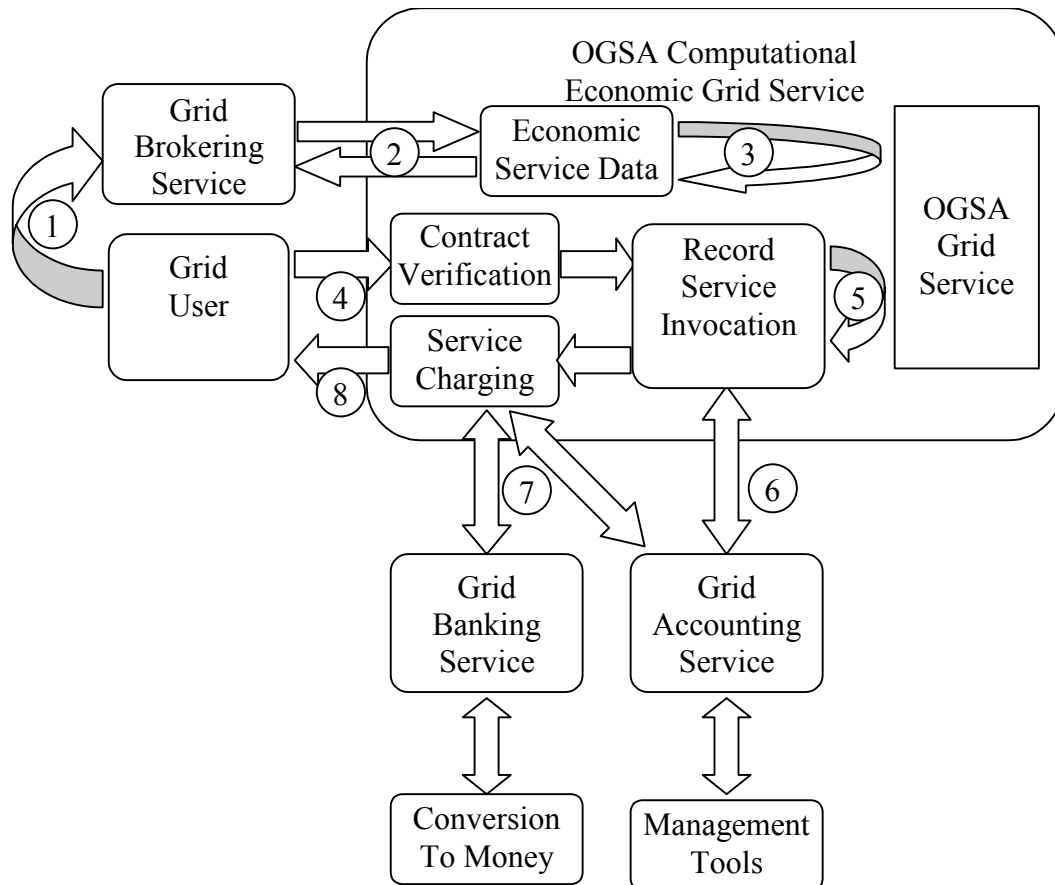


Figure 1: Proposed OGSA 'Strawman' Computational Market Framework

Work Package 3: Grid Accounting Service

The Grid Accounting Service interface defined in WP 2 will be implemented to provide a reference implementation for the Grid community. The service, with appropriate client side tools, will enable resource owners to monitor usage and utilisation of their Grid resources. This service will be OGSA compliant enabling its contribution and deployment to the next release of the Globus Toolkit (GT3). There will be no restriction on the 'types' of resource utilisation that can be recorded (e.g. memory, CPU time, etc.) and then accounted for by the associated management tools. It is recognised that there are privacy issues relating to this accounting information and access will need to be controlled within a defined public policy.

Deliverables

D3.1: Development of service implementation to record resource usage.

D3.2: Development of an API to allow service providers to record accounting information in the accounting database.

D3.3: Development of client tools to enable users and resource providers to extract stored accounting information.

Work Package 4: Grid Banking Service

An OGSA compliant prototype Grid Banking service will be implemented to provide a secure charging and payment mechanism for resource usage. We will investigate existing commercial electronic payment methods and their compatibility with this service together with mechanisms for secure payment. Following a survey of these systems we will re-implement the banking service to leverage one of these existing commercial systems to support wider production deployment.

Deliverables

D4.1: Implementation of a Grid Banking service to enable exploration of resource trading and the verification of the service interface.

D4.2: Survey of existing and emerging electronic payment mechanisms.

D4.3: Implementation of a Grid Banking service that leverages an existing and established electronic payment mechanism.

Work Package 5: Grid Brokering Service

WP 3 & 4 provide the basic mechanisms for building a computational economy: recording service usage and the ability to receive payment for the use of a service. The matching of resource requests by a client to the capabilities and cost of a service (this work package) uses the meta-data ontologies specified in WP2. The decoupling and separation of the mechanism used to set the price from the process used to match the service capability to the consumers requirements is essential as it allows sophisticated negotiation mechanisms to be introduced at later stages of the project or to be contributed by other groups.

The resource brokering service will accept work plans from the consumer and obtain predicted costs for this work from the available service providers. This process will be recursive allowing a costed workplan to be generated directly by the service provider or indirectly by another resource broker acting on behalf of the service. The ability to recursively request resources allows a resource broker to modify the cost of the returned workplans allowing it to either make a charge for the service, or to attract business by adopting a loss-leading strategy. Allowing third-party brokers to profit will enable the development of increasingly sophisticated brokers. The resource broker may finally return all the costed workplans to the consumer, or alternatively process the workplans in accordance with a policy (i.e. Quality of Service, QoS, constraints) supplied by the consumer. This may specify acceptable cost, acceptable turnaround, a fitness function, or a limit on the number of offers returned to the consumer. A consumer is free to pass on as little or as much of its service selection policy to the brokering service and to reject any offers made by the broker.

The costed workplan will form the basis of the service contract. The system is, of course, free to implement the contract in anyway it sees fit (see Workpackage 8). Mechanisms to deal with broken contracts will also be required (for example, by the inclusion of penalty clauses in the contract).

The focus of this workpackage is essentially concerned with pre-deployment issues for Grid applications. The costing of workplans will entail using performance data relating to, possibly multiple, implementations of the components which will make up an application. This data may be prediction data or data obtained from previous executions of the component, for example. Workpackage 8 addresses run-time issues for Grid applications.

Deliverables:

D5.1: Implementation of a brokering service that supports recursion to other brokers and Grid services.

D5.2: Extend the basic brokering to include policy driven cost modification.

D5.3: How to use performance data in costing workplans and developing pre-deployment contracts.

Work Package 6: Exemplar e-Science Applications

Geodise

The Southampton e-Science Centre is leading the Geodise EPSRC Pilot project. This aims to develop a Grid-based platform for design search and optimisation involving a variety of sophisticated computational services: parallel hardware, numerical solvers and databases. Southampton and Swansea propose to use this as an exemplar to test the requirements for a computational market. They will “deconstruct” the application into its component services and explore how these could be provided and the project reconstructed by buying in the appropriate services using the market mechanisms and protocols identified by this project. The key knowledge will be the level of detail required to specify the required services adequately.

Engineering design search and optimisation is the process whereby engineering modelling and analysis are exploited to yield improved designs. The key to successful optimisation lies in maximising the improvement to the design given the available computational resources. It is thus necessary to determine the cost of performing each analysis, establish what resources may be available at what price, devise a suitable optimisation schedule on the basis of the available resources, and then allocate jobs to these resources. The Geodise project is focussing on grid-enabled optimisation for Computational Fluid Dynamics (CFD) and a key factor in determining the computational cost of the calculations is the size of the mesh to be employed for the calculation. The University of Wales, Swansea, will provide a grid-enabled meshing service which has a well-characterised performance-cost model and this will be used to provide ‘third party’ meshes for the Geodise CFD optimisation calculations. This combination of component services from various providers will provide a key demonstration of the application of the brokering methods developed in the previous work packages. The overall project will provide a significant enhancement to the current generation of grid-middleware, which does not provide much of the functionality essential to grid-enabled tasks such as optimisation and design search.

Deliverables:

D6.1: Development, Integration and deployment of brokering service into Geodise Computational Service.

D6.2: Grid-enabling meshing software.

D6.3: Integration of 6.2 into brokering service (and Geodise).

eSTAR

The Astrophysics Research Institute of Liverpool John Moores University is leading the eStar e-science demonstrator. This aims to develop and test the infrastructure necessary for the deployment and optimal use of a global network of robotically controlled telescopes and data archives accessed by Intelligent Agents that act as the astronomer’s proxy. One key issue is the allocation and tracking of the scientifically and financially valuable telescope time across the heterogeneous network of telescopes and archives. Data from such facilities has a dynamically changing value depending not only on the inherent data quality due to the size and design of the originating instrument, but also the prevailing environmental (weather, atmospheric stability, sky brightness) conditions. In addition for some projects there are strong timeliness constraints where data may have considerable value if

obtained in the next ten minutes but none if obtained after that period. The current model within eSTAR abstracts some of these quantities into a single value (ignoring the rest altogether) and makes no attempt at all to implement any accounting of resources. ARI will recast this current (very basic) approach into the market model in order to confirm that the derived architecture and services from WP's 2 – 5 are well fit to the actual problem, and implement the new approach within the existing demonstrator to provide a system for trial deployment.

Deliverables:

D6.4: Integration of e-STAR with the dynamic resource brokering model.

Work Package 7: Trial Deployments

We will deploy and test the prototype implementation produced in WP 3, 4 & 5 over critical commercial and scientific Grids.

CSC will deploy the new tools/standards for one or two of its reference customers, where it aligns with their lifecycle of IT refreshment. Initially this will be for internal use only. Subsequently they would connect two or more such Grids, including their e-HPC facilities, and then to roll out full accounting and invoicing, having by then ensured these standard processes meet CSC's needs.

The second trial will be over the UK e-Science Grid. The Grid Support Centre will lead a project where selected regional centres, almost certainly including those in this project will deploy the enhanced middleware and conduct experiments in the support of key e-Science applications, including the Southampton exemplar, and resource sharing.

This workpackage will also examine the interaction of the UK Grid economy with other Grid economies via small-scale experiments and simulation in trading resources. In this way, the feasibility of trans-national Grid economies can be explored independently of the political and administrative considerations of transfer of cost resources across organisational and national Grid domains (since the resource transfer in the experiments will be negligible).

A market enhanced eSTAR trial deployment will be conducted to overseas telescope sites run by UK organisations (but which also have non-UK users) in Spain and the US, providing a test of a transnational Grid economy within a complex scientific community

Deliverables

D7.1: Deployment of an economic trading grid within CSC managed resources.

D7.2: Deployment of an economic trading grid within the UK e-Science Grid.

D7.3: Interaction between two discrete grids through electronic trading schemes.

D7.4: Deployment in an international astronomy community.

D7.5: Combined deployment of Geodise + meshing service with dynamic resource brokering.

Work Package 8: Negotiating Mechanisms within Computational Economics

Once the basic infrastructure for recording service use, finding a price for a service and the supporting payment mechanisms have been provided by the early deliverables of WP 3, 4 & 5 we will be able to prototype increasingly sophisticated charging regimes to optimise resource utilisation. We propose to develop a price negotiation interface that can be invoked by the resource broker regardless of the trading mechanism. This 'pluggable' approach to the negotiation and trading mechanism will allow us

to explore different trading mechanisms, static and dynamic pricing, contract adherence and performance based application mapping.

Deliverables

D8.1: Static Price Trading

This is a simple well established trading mechanism where the cost of service invocation is fixed for a particular user (or group) at a particular time for a specific 'quality of service' (e.g. peak, off-peak, background with potential pre-emption, priority etc.). A serviceData schema to support static price trading will be developed and used to explore this mechanism.

D8.2: Dynamic pricing through auctioning

We will exploit the infrastructure developed in earlier work packages to implement and explore a pricing strategies based upon several standard auctioning techniques (e.g. English, Dutch, Hybrid and Sealed Bid auctions) in order to find globally optimal application mappings.

D8.3: Contract compliance.

All service usage thorough the economic model will be expressed as a service contract. The delivery of the service contract will be monitored through the information collected by the accounting service. The service provider may also wish to monitor applications to ensure that they stay within the agreed contract (e.g. do not consume resources in excess of what has been agreed). Likewise, the consumer will wish to ensure that the service provider will be able to deliver on any negotiated contract. Work on this is already in progress under the GrADS project in the US and an important component of this Workpackage will be to establish a complementary approach to this work.

D8.4: Performance Engineering

We will investigate mechanisms for using performance and cost information to obtain the most cost-effective means of executing applications. Imperial, Manchester and Southampton have considerable expertise and experience in the use of application performance models in guiding mapping and Imperial and Manchester are addressing these issues in the EPSRC Pilot Reality Grid. Two aspects of this will be considered: *pre-deployment performance analysis*, where an application is constructed in such a way that its performance characteristics are known and can be used to inform selection of the appropriate resources (this is the subject of Workpackage 5), and *performance steering* where the dynamic behaviour of the resources and the application are monitored during execution and used to guide resource selection. This workpackage therefore addresses post-deployment issues for Grid applications.

For example, a user who wishes to run a Gaussian analysis task with constraints limiting the latest completion time and maximum cost will need to provide information relating to the application's performance on various resources. This information may be extracted directly from the application as it executes, from meta-data provided by the developer or information extracted from the accounting information relating to previous runs. The requirements of the user (e.g. that Gaussian is installed) may eliminate certain resource providers. Once executing, components of the application may be redeployed on additional resources as they become available allowing the "system" to meet the contract more efficiently (at greater profit to the service provider). The Performance Engineering deliverable will address techniques and mechanisms for gathering and utilising run-time performance information and catalogue techniques for making components "malleable", (that is, so that they are able to alter their performance behaviour as circumstances change). A simple example is the ability of a parallel component to respond to a request to continue execution on a different number of processors. A performance steering architecture will be developed. The use of such mechanisms to support the higher level policies imposed by computational economies will also be investigated..

Work Package 9: Dissemination

Information relating to the project will be disseminated through a public web site that will maintain a copy of the protocol specification and versions of the reference implementation as they become available. The global nature of Computational Grids mean that it is essential that there be widespread agreement on the protocols and services underpinning this project. We will therefore develop our

4. Partners' Background and Roles

4.1. CSC's involvement for use of Grid tools in Commercial/Industry arenas

<http://www.csc.com/>

Key Individuals

John Rawlins, Nigel Healy, Gerry Todd.

Background

Computer Sciences Corporation (CSC) is a large IT Services company with a global presence.

CSC manages and delivers the IT Services to its clients across a broad range of industries, for both government and industry.

CSC is the lead contractor for managing many of the world's largest HPC sites (incl. NASA Ames, NASA Goddard, Wright Patterson Air Force Base, Vicksburg Navy base, BAE SYSTEMS and the CSAR Service). These HPC sites deploy large scale, modern HPC technologies from IBM, HP/Compaq, SGI, Cray and Sun. CSC also manages private networks for its clients as well as providing mainframe and web hosting services, customer service facilities, financial applications and IT consultancy. CSC, in particular, manages over 100,000 desktops and servers.

CSC is keen to support emerging Grid standards and in particular to help our clients to deploy Grids for internal and external consumption. By connecting supply and demand, the Grid is most likely to become a commercial success.

CSC proposes to support and to use the Grid tools in real world commercial scenarios through a progressive phased deployment plan, so that CSC can influence the Grid tools, ensuring they meet CSC's and, most importantly, our clients' requirements, with appropriate interoperability between IT technologies.

Scope of CSC's Involvement

1. Consultancy. We propose to work with the others in this research activity to help derive new industry standards that will meet CSC's medium/long term requirements, in particular to ensure that suitable resource accounting, scheduling/brokering, security and system management functionality is in place.
2. Alpha Test. We propose to deploy the new tools/standards for one or two reference customers, where it aligns with their lifecycle of IT refreshment. Initially this will be for internal use only. Subsequently the plan would be to connect two or more such Grids, including CSC's e-HPC facilities, and then to roll out full accounting and invoicing, having by 1) ensured these standard processes will meet CSC's needs.
3. Support agreement for new standards. Participate in the GGF.
4. Demonstrate the working solution to interested parties. We would anticipate maintaining operational systems which may be demonstrated to those seeking to implement or use Grid components and as a vehicle to test layered products.
5. Beta Test. We propose to develop and maintain a test harness through which Grid product releases can be proven before installation and release for production working.

Estimate of cost to CSC's involvement

CSC estimates our support will average two full time equivalents (FTE), which is equivalent to a CSC investment of £150K p.a. inclusive of overhead charges and HPC machine time for development purposes.

4.2. GridX's Involvement

Key Individuals

Dr Paul Calleja, Mr Terry Benga

Background

GridX is a new start-up HPC/Grid company currently entering the final stages of venture capital fund raising round (in the order of £5M), due for completion by September 2002 when the company will become operational.

GridX will provide leading edge commodity compute, storage and visualisation solutions through the provision of both hardware and software to the traditional HPC markets plus emerging and rapidly growing enterprise markets through the development and deployment of Grid middleware.

In addition GridX will be fully engaged with the UK e-science programme by acting as a knowledge transfer mechanism from the e-science centres and projects to industry and commerce. GridX will develop a commercial strength Globus/OGSA environment by implementing additional functionality where necessary. Key areas of development will be cluster management software, scheduling, security models and utility computing functionality.

We see Grid commercialisation evolving through 3 phases. The initial phase will be based around consultancy and Grid deployment for intra-grids for increased utilisation of pre-existing resources. Secondly, offloading peak demand and non-critical or sensitive enterprise application via utility Grid middleware to off site data-centres. Thirdly, more widespread and pervasive adoption of Utility compute and storage models, with emergence of new working practices and business models.

GridX builds on a strong knowledge base in HPC and Grid computing. Dr Paul Calleja was founder and head of Compusys HPC for the last 3 years, during this time Compusys HPC rose to be one of the leading commodity HPC companies in the U.K. installing key machines at RAL, Imperial, Southampton, Bristol, and the AWE to name but a few, with active involvement with grid projects at Southampton and Imperial.

GridX project contribution

- 1) **Consultancy.** GridX has completed an in-depth feasibility study of utility computing business models for both the commercial and academic users with a mixed pricing scheme based on differential customer SLA's. In addition to the Utility business case GridX is also actively investigating vital enabling software components which are needed for a fully functional utility computing solution stack, from data-centre through broker to end-user. This expertise will be contributed to WP 1 & 2.
- 2) **Co-development.** GridX will co-develop utility OGSA tools and implementations with Imperial by sponsoring one full time FTE based at the College and contributing a further developer's time from within the company. This will open up our in house R&D and grid reference site resources. This will contribute to WP 3,4 & 5.
- 3) **Testing.** GridX will deploy utility computational tools within our OGSA distribution development sites, which consist of both academic and commercial early adopters providing feedback within WP7.

- 4) **Deployment and support.** GridX will deploy and support utility grid solution stacks in commercial and academic environments and foster an active computational and storage free market.
- 5) **Key Collaborations.** GridX has key collaborations with Intel, Dell, Quadrics and QinetiQ and these collaborations will greatly benefit the project as we widen the scope and outreach of the work.

Estimate of GridX's contribution

We estimate the value of this contribution at £115K p.a. Based on 2 FTE's @ £50K p.a. & 10% of a Director's salary.

4.3. NAG's Involvement

<http://www.nag.com/>

Key Individuals

Brian Ford, Andrew Giles, Mike Dewar, Steve Hague

The Numerical Algorithms Group (NAG) Limited is an SME (currently around 100 employees world-wide) that has been an international leader in the development, distribution and support of high quality mathematical software since its inception as a company in 1976 (and as a collaborative project in 1970). Its principal product family is a large, tailored collection of numerical and statistical algorithms (the NAG Library). These algorithms, in several different language forms, are extensively used in a very wide range of applications in scientific, technical and analytic computation. Traditionally, NAG's primary audience has been end users (researchers in universities and industry, for example, who would write their own programs that would then call NAG Library software) but an increasing proportion of NAG's business now comes from licensing of its numerical and statistical library components in the computational packages of other software vendors (thus, more B2B than B2C in e-commerce parlance). In terms of vertical markets, the most quickly growing sectors for NAG are in the banking and financial services area, in which organisations are increasingly turning to sophisticated analytical techniques (including data mining and cleaning, forecasting and optimisation) to gain a competitive advantage.

NAG's numerous clients and partners include all the major hardware vendors, PeopleSoft, OriginLabs, Merrill Lynch, Waterloo Maple, Zeneca, Philips, DERA, Mercedes, BAe, Daimler Chrysler, the National Physical Laboratory and various other government research establishments, and universities around the world.

Other products in the NAG portfolio include programming support tools and an advanced, data visualisation and application building system, IRIS Explorer. That latter product is particularly suited to remote computing because of its support for cooperative visualisation and computational steering (thus allowing research collaborators, or perhaps client and remote service provider in a commercial context to act on a joint basis to monitor, or indeed even alter, the course of a protracted computation).

Scope of NAG's Involvement

NAG is following closely and participating in Grid and e-Science related activities through its various national and international contacts, and in particular it is associated with the recently confirmed e-Science Open Call project, Visualisation Middleware, to be led by the University of Leeds. It is also the coordinator of a European Union FP5 Semantic Web project, MONET (Mathematics on the Web) that has a direct and complimentary bearing on the programme of work of this proposed project.

NAG shares the enthusiasm and willingness of the other industrial partners to be involved in the proposed project because the Computational Grid is already a reality in the research sphere and may well become a pervasive commercial reality within the foreseeable future. If that does happen, it is

essential for any forward-looking IT company to anticipate rather than react. Early exposure to and trial experience of the issues identified in this proposal are therefore very apposite and timely for NAG.

We anticipate that our involvement will mainly take the form of sharing knowledge, both technically and commercially, with our project partners, in Work Packages 1 to 5 and 8 in particular, and we will take a keen watching brief in WPs 6 and 7. We will offer practical assistance to partners where it would be useful and practical to do so, and we also anticipate playing an active role in WP9 (Dissemination). Our main line of interest will be in issues that are of particular significance in the areas of:

- complex computation (e.g. how to choose by automated means, the right computational method, or how to judge whether a computation has succeeded or failed, and in the latter case, who pays for a re-run!) and
- component supply (Use case 1.5) – what will be the additional requirements placed by a commercial Grid service on NAG (which has in fact been a component supplier since its inception)?

We will ensure that there is effective information flow between this project and others in which we are involved, specifically the EU-funded MONET Semantic Web project. <http://monet.nag.co.uk/>

Subject to further discussion and more detailed planning,, we are willing to make available licences for use of NAG products (or appropriate parts thereof) to serve in a guinea-pig capacity in experimental services.

Estimate of Cost of NAG Involvement

The nature of our envisaged involvement will require the participation, sometimes separately and at other times jointly, of senior technical and commercial staff. We expect that this involvement might average of the order of 1.5 days per month. Given the relative seniority of this involvement, and applying standard full commercial overheads, this equates to a contribution in the region of £20,000 per year. Subject to confirmation of the details of the specific licences that may be required, we estimate that the notional commercial value of NAG product software licenses granted would be in the region of £8000 per year. Thus the combined estimated annual contribution is around £28,000.

4.4. Real Time Engineering's Involvement

<http://www.rtel.com/>

Key Individuals

Jim Henderson, Conway Williams, Alastair O'Brien, Alastair Deacon.

Background

Real Time Engineering Limited (RTEL) is a medium sized enterprise (~120 staff of whom around 100 are professionally qualified in either computer science or technologies specific to our customer base). We have been trading since 1988 and have offices in Glasgow, Livingston, Aberdeen and at Heathrow Airport. Our main clients are both national and multi-national organisations and we have carried out consultancy and supply work in most major countries.

RTEL provide specialist software, requirements capture/design consultancy and system integration services to a range of industries whose prime criteria is for high integrity, fast response systems. Our customer base includes all the major oil companies, National Air Traffic Services & Eurocontrol, government regulatory bodies (e.g. OFGEM, CAA), airport operators, (e.g. BAA), the energy industry (e.g. British Energy, Centrica, SSE, Innogy) and major financial/insurance institutions (e.g. RBoS, Intelligent Finance, Abbey National).

Our products include energy and gas trading software (RT-X), designed to offer an end-to-end solution to energy traders in the New Electricity Trading Arrangements (NETA) environment, gas trading facilities for the deregulated market and share trading and insurance trading for IFA's. We have pioneered web-enabling of industrial and legacy systems for most of our major customers and our latest mainstream offering (Airport Roster Management) is implemented in a distributed Web Services environment using the latest web sharing software techniques.

RTEL would be delighted to offer our expertise in designing and implementing both trading systems and web services to the Market for Computational Services project.

Scope of RTEL's Involvement

1. Consultancy. We will provide expert advice to the research activity in the areas of setting up a practical trading framework leading to the evolution of a viable computational marketplace during work packages (WP) 1 to 5.
2. Practical Assistance. We will furnish practical assistance during the prototype stages of the project and assist in the development and deployment of the functionality during WP 6 and 7. This may well include access to Intellectual Property used within existing trading products.
3. Commercial Advice. We will use the benefit of our experience in a wide range of trading environments to assist in the development of the more advanced price trading and contractual services models for WP 8.

Estimate of Cost of RTEL's Involvement

Since our involvement in the project would be of a strategic advisory nature (at least initially) we have proposed to provide advice from a number of experienced senior staff during the project timeframe. We would estimate the involvement to be at the level of roughly 0.5 full time equivalent (80% Principal Consultant & 20% Director level). At standard commercial rates this represents a cost of £125K p.a.

4.5. Sun Microsystem's Involvement

<http://www.sun.com/>

Key Individuals

Peter Overton, Sun Microsystems UK

Wolfgang Gentzsch, Sun Microsystems Sun Inc

VISION: The Network Is The Computer™

Sun was founded with one driving vision. A vision of computers that talk to each other no matter who built them. A vision in which technology works for you, not the other way around. While others protected proprietary, stand-alone architectures, Sun focused on taking companies into the network age,

providing systems and software with the scalability and reliability needed to drive the electronic marketplace.

Market Requirement

Take an enterprise grid for a financial institution in Europe, for example, which provides portfolio and risk management services to its wealthy customers and investors. In the past, different departments have developed different, very market-specific solutions, each within their own home-grown departmental HPC environment. From an enterprise perspective, this solution is very inefficient. There is replication, system administration, maintenance, tuning, training, etcetera, in every department. A new enterprise grid solution provides the potential economies of scale, access to one common HPC service for all departments, reliability and quality of services, reduced hardware and software costs, reduced operational cost, and increased productivity.

Research requirement.

Sun Microsystems currently offers two ways of combining its two software stacks, Sun Grid and Sun ONE. The first solution is via the industry-proven iPlanet Portal Server for secure, remote and transparent access to Sun ONE Web services and to Sun Grid computing services. This obviously leads to one global access point for businesses, which first want to process data in their grid, and make the results available as a service to employees or customers. The second solution is via Sun ONE Connectors to major applications, like ERP and CRM, or simply via JSP or EJB Containers that wrap applications inside a Java Bean. Once integrated with our Java Beans technology, the grid application is easily integrated into a Web service. This technology allows our customers to build their own scalable compute services and offer them through Sun ONE, as services on demand. In such a combined environment, Sun ONE manages users, customers, communities, policies, and even access to grid services, while the grid manages computing, data processing, storage, efficient utilisation of resources, collaboration, and more.

Background:

Building upon the Service Driven Network Sun has been talking for several years about the Service Driven Network vision. That vision is based on a responsive digital network that gives customers "'dial tone' reliability while delivering new services that reside on the network." In essence, it's bringing companies to the best quality of service they can expect from a network, at a cost that is still practical for running a business.

The Services on Demand Vision

The network is all about servicing your communities, stockholders, customers, and employees. Increasingly, these services are delivered to communities via software functionality that has migrated to the network. That software may be in distributed form, saved in pieces in directories on the network, or called upon to perform specific functions. It may involve Web applications, but may also embrace different innovations in things like Web services or other component software models. Sun includes all those software methods in its Services on Demand vision. Services on Demand is the point where information assets are fully leveraged for business benefit in the form of services -- using any technological way that's cost-effective and suitable for a company, be it Web applications provided as services, or components of applications stored in central directories assembled in real-time to provide services. It's all about the ability to use all these software approaches, according to the phase a company is in, and the timing that's right for that business.

Evolution, Not Revolution

In fact, in Sun's view of Services on Demand, software technologies don't disappear, they evolve to include legacy as well as new technologies. These technologies build off of each other and intermingle, so that customers don't have to start from scratch. It is Sun's view that software evolution, not revolution, is the best approach for maintaining ROI and building effective business practices.

TARGET AUDIENCE:

Every network environment is a candidate for Grid Computing, but the product is especially attractive to Electronic Design Automation, Life Sciences, Scientific Research and MCAE industries. Other industries are Financial Modelling, Chemical Engineering, Oil and Gas and there is interest from leading edge commercial applications for business modelling.

It is proposed that Sun Microsystems collaborates with Imperial College and the Market for Computational Services Project in the further development of the Java, Solaris, Sun Grid Engine and SunONE frameworks, architectures, technologies and product implementation, to enable a Sun reference platform for OGSA. It is envisaged that this reference platform would manifest itself in open source and Sun Product implementation and a Reference Implementation (RI) of a 'Chargeable Grid Service' and a 'Resource Usage Service' interfaces that will be developed during the project.

Estimate of cost of Sun's involvement

Sun Microsystems estimates our support will average two full time equivalents (FTE), which at the standard rate for our customers would be £150K p.a. This level of support is at the discretion of Sun Microsystems, and subject to agreement of research project objectives and delivery plan.

4.6. SGI's involvement for use of Grid tools in Commercial/Industry arenas

<http://www.sgi.com/>

Key Individuals

Andrew Grant, Martyn Foster.

Background

Celebrating its 20th year, SGI, also known as Silicon Graphics, Inc., is the world's leader in high-performance computing, visualisation and the management of complex data. SGI products, services and solutions enable its technical and creative customers to gain strategic and competitive advantages in their core businesses. Whether being used to design and build safer cars and aeroplanes, discover new medications and oil reserves, predict the weather, entertain us with thrilling movie special effects or provide mission-critical support for government and defence, SGI systems and expertise are empowering a world of innovation and discovery. The company, located on the Web at <http://www.sgi.com>, is headquartered in Mountain View, California, and has offices worldwide.

SGI provides capability to the world's most demanding grid environments. From the Information Power Grid distributed across the multiple NASA sites in the US, to HPC on demand facilities based at Manchester University, SGI is leading the way with solutions that tackle the real and diverse requirements for grid computing in a multitude of technical compute scenarios. SGI is unique in its market position of supplying systems and integrated solutions from desktop workstations through to the world's most advanced supercomputing resources, all integrated with high performance storage and visualisation components.

SGI is taking a leading role in the evolution and adoption of grid computing, complementing existing work with developments in Visual Area Networking, Distributed Filesystems, and performance management. SGI also provides accounting tools via the SGI Comprehensive System Accounting system used in some of the world's largest HPC environments. SGI delivers grid solutions to its customers via a global Professional Services operation and via a strategic alliance with Platform Computing to deliver supported off the shelf middleware solutions.

SGI is keen to enable new models of grid computing, via the integration of accounting and billing technology under common industry standard interfaces to enable our customers to operate compute on demand facilities with minimal difficulty maximising their investment in SGI systems and resources.

Scope of SGI's Involvement

1. Consultancy. SGI will work with the partners in this proposal to provide consultancy and steering to the program. Drawing on the large number of our customers in Science and Industry who have demanding accounting and auditing requirements, we will provide input to define the baseline requirement for accountancy, audit, and billing needed to achieve market acceptance of any protocol and implementation in the space defined by this proposal.
2. Test and Review. SGI will test and assess implementations against the needs of several significant organisations who are interested in migrating to grid based architecture.
3. Support for standards acceptance. SGI is a platinum sponsor of the Global Grid Forum, and has working relationships with the developers and commercial distributors of Globus. SGI will support the proposals made by this project to ensure that both visibility and acceptance of the work is present, with the end goal of APIs and interfaces forming industry standard mechanisms accepted by GGF.
4. Technology. SGI will supply a key technology piece in the form of SGI CSA technology described above, that will be invaluable to the project. CSA gathers and records accounting data, generating flexible reports and billing guidance. SGI will supply this technology and source code together with "Process Aggregate" mechanisms that combine distributed jobs into a single accounting container.
5. SGI will provide a configured development vehicle (workstation) configured with Platform LSF, SGI CSA and SGI PAGG, for testing and research. SGI will also provide access to a range of facilities in order to test components of the developing system against resources that may not be available to the partner sites but must be part of any accounting system (large shared memory architecture, high performance graphics pipes, SAN based storage, real time systems etc).

Estimate of cost to SGI's involvement

SGI estimates our support will average about one full time equivalent (FTE), which is equivalent to an SGI investment of £75K p.a., i.e. a total of £150K, inclusive of overhead charges and HPC machine time for development purposes.

4.7. Oracle

<http://www.oracle.com/>

Key Individuals

Colin Brown, David Pearson

Background

Oracle Corporation is the world's largest enterprise software company, providing enterprise software to the world's largest and most successful businesses. With annual revenues of more than \$10.8 billion, the company offers database, tools and application products, along with related consulting, education, and support services, Oracle was the first software company to develop and deploy 100 percent Internet-enabled enterprise software across its entire product line: database, server, enterprise business applications, and application development, and decision support tools.

Scope of Oracles Involvement

Oracle is keen to have a relationship with the Computational Markets project, particularly to help explore the business model and mechanisms for the provision of data services, and is willing to assist the project by performing reviews and in attending appropriate workshops.

In addition Oracle sees the need to link the OGSA-DAI and the Computational Markets projects to exchange information and ideas and David Pearson has expressed a willingness to assist in this process

Estimate of Cost of Oracle's Involvement

Oracle is already committed to making a substantial contribution to the OGSA-DAI programme of work. For this reason it is not able to make a formal contribution to the Computational Markets project which would attract matched funding from the DTI.

5. Regional Centres and Universities Involved

- London e-Science Centre (John Darlington, Peter Harrison, Steven Newhouse, Berc Rustem) Grid middleware, Software Component technology, Performance engineering, Stochastic optimisation, Computational economics. <http://www.lesc.ic.ac.uk/>
- CLRC e-Science Centre and Grid Support Group (David Boyd, Rob Allen) Grid middleware, Grid administration and policy, Web Services.
- North West e-Science Centre (John Brooke, John Gurd, Jon Maclaren, Stephen Pickles, Graham Riley) Grid Middleware, Component technology, Performance engineering, Resource brokerage. <http://www.esnw.ac.uk/>
- Liverpool John Moores University, Astrophysics Research Institute (Iain Steele) Astronomical instrumentation and software. <http://www.astro.livjm.ac.uk/>
- Southampton e-Science Centre (Simon Cox, Nigel Shadbolt, Nick Jennings, David de Roure) Web services, Performance engineering, Agent technology and negotiation. <http://www.e-science.soton.ac.uk/>
- University of Wales at Swansea (Nigel Weatherill) Numerical modelling, Mesh generation. <http://www.swan.ac.uk/civeng/>

5.1. Astrophysics Research Institute Background

Key Individuals

Dr Iain Steele, Dr David Carter, Mr Jason Etherton

Background

The Astrophysics Research Institute (ARI) of Liverpool John Moores University is a medium size (40 staff & students) astrophysics group with particular specialisations in astronomical instrumentation and software. It is currently building the world's largest robotic telescope (the 2.0-metre Liverpool Telescope) sited on La Palma in the Canary Islands through the university spin off company Telescope Technologies Ltd, and undertakes contract research for other telescope users such as the Faulkes Telescope Corporation and well as carrying out an academic research programme funded through HEFCE and PPARC. The ARI software group is lead by Dr Iain Steele, and employs 4 full time software engineers. As well as general software engineering skills in Java, C and hardware control, the group has particular e-science expertise in Globus, LDAP, XML and network programming.

In collaboration with the Physics & Astronomy Department of the University of Exeter, the ARI has successfully delivered an e-science demonstrator eSTAR (e-science Telescopes for Astronomical Research) that uses the Globus grid tools to implement the infrastructure for a global network of robotically controlled telescopes and data archives accessed by Intelligent Agents that act as the astronomers proxy.

In order to move eSTAR from a demonstrator to a system that can be deployed on professional telescopes however, a number of issues remain to be addressed. One key issue is resource allocation and tracking. The computational market approach offers an excellent opportunity to address these problems, and one can eventually envisage all telescope time being managed in such a fashion.

Scope of ARI Involvement

WP1: ARI would provide input to, and participate in the development of, the use cases D1.3 (Scientific User), DN1.5 (Storage Service Provider) and DN1.6 (Instrument Service Provider). This would be driven from the perspective of astronomical instrumentation, but we would also consult with colleagues in other physical sciences to bring their requirements into the process.

WP2: ARI would participate here in providing advice and testing against the use cases derived in WP1.

WP6: ARI would re-engineer their current eSTAR system to incorporate the market model, adding support for at least the grid accounting and brokering services, and possibly the grid banking service.

WP7: ARI would test the re-engineered eSTAR systems on UK and intentional networks in order to prove the computational market concept, especially in use cases D1.3 and D1.6.

Appendix A: Allocation of Deliverables

Deliverable	Lead Site	Contributors	
		Academic	Industrial
D1.1: Application Service Provider Use Case	Soton	Swan	NAG, SGI
D1.2: Computational Service Provider Use Case	UoM	IC	CSC
D1.3: Scientific User Use Case	Soton	Swan	NAG
D1.4: Resource Reseller Use Case	UoM	IC	RTEL
D1.5: Component Supplier Use Case	IC	Soton	NAG
D1.6: Instrument Provider Use Case	ARI	UoM	
D1.7: Utility Computing Use Case	IC	Soton	GridX
D1.8: Data Provider Use Case	IC		Oracle
D2.1: Resource Schema	IC	UoM, ARI	RTEL
D2.2: Trading Schema	UoM	IC	RTEL
D2.3: Accounting Service Definition	UoM		CSC
D2.4: Banking Service Definition	IC		NAG
D2.5: Brokering Service Definition	UoM	IC	RTEL
D2.6: Chargeable Grid Service	IC		GridX, CSC
D3.1: Record Service Use	UoM		CSC
D3.2: Management API	IC		CSC
D3.3: Client Accounting Tools	Soton		NAG, SGI
D4.1: Basic Grid Banking Service	IC		
D4.2: Survey of Electronic Payment Systems			RTEL
D4.3: Real Banking Service	IC		
D5.1: Recursive Brokering Service	UoM		CSC
D5.2: Policy Driven Brokering	IC		GridX
D5.3: Use of Performance Data in Resource Pricing	UoM	IC	
D6.1: Integrate Brokering into Geodise	Soton		
D6.2: Grid enable meshing software	Swan		
D6.3: Integrate Meshing Software into Brokering service	Swan		
D6.4: Integrate eStar Application with Brokering	ARI		
D7.1: Deployment within CSC	CSC		
D7.2: Deployment within UK e-Science Grid	GSC	IC, UoM, Soton	
D7.3: Interaction between two grids	GSC		
D7.4: Deployment in international astronomy testbed	ARI		
D7.5: Deploy integrated Geodise and Meshing service	Soton		
D8.1: Static Price Trading	IC	UoM	
D8.2: Dynamic Auctioning	UoM	IC	
D8.3: Contract compliance	Soton	UoM	
D8.4: Performance Engineering	UoM	Soton	
D9.1: Project Web Site	IC	All	
D9.2: Engagement with GGF and development of WG activities	IC	All	Sun, SGI
D10.1: Project Management	IC	All	All