

Core Grid Middleware

The North-East Regional e-Science Centre is involved in a large number of research projects that rely on the design and development of a Grid-based infrastructure. Building a large number of different infrastructures would be time-consuming, difficult and risky; currently available Grid middleware is relatively immature and subject to frequent change, while relevant knowledge and experience among e-Science researchers is understandably very limited. A further issue is that many e-Science projects—both at Newcastle and elsewhere—are currently based on Web Services, with the intention of moving to Grid Services at some point in the future. The method of making the transition is not yet clear, nor is the extent of the extra work, delays and risks this will introduce into projects.

Consequently, it was decided to analyse the requirements of all the Newcastle projects and to design and build a common *Core Grid Middleware* package that, in its initial release, will incorporate four key Grid Services (Figure 1).

Following that initial release, the Core Grid Middleware will continue to evolve by incorporating the “best-of-breed” OGSA Services produced by the Grid community or, if deemed necessary, developing them in house.

The experience gained from the development and deployment of the Core Grid Middleware will be valuable for the UK e-Science community wishing to adopt the Open Grid Services Architecture; for any new Grid Services related research projects; and as feedback to the process of building the UK Grid Services-based infrastructure and the Open Middleware Infrastructure Institute.

OGSA-DAI

Many of the research projects in which the NEReSC is involved require access to database management systems over the Grid. The Open




Figure 1. Grid Services in the initial release of the Core Grid Middleware

Grid Services Architecture – Database Access and Integration (OGSA-DAI) service provides a consistent way to access relational and XML data on the Grid. The OGSA-DAI implementation is included as is in the Core Grid Middleware.

OGSA-DAI DQP

NEReSC is directly involved in the design and implementation of the OGSA-DAI Distributed Query Processing (DQP) Grid Service (reference to your section on OGSA-DAI DQP), which enables Grid applications to run queries on distributed data resources. A number of e-Science research projects will greatly benefit from the inclusion of OGSA-DAI DQP in the initial release of the Core Grid Middleware.

Workflow

The ability to capture and enact computations is important for all the e-Science projects. Existing work on workflow execution in the ^{my}Grid project is the basis for the Workflow Enactment Grid Service (WEGS) for the Core Grid Middleware. Currently, the ^{my}Grid enactment engine is based on Web Services and so

NEReSC is producing WEGS to be OGSi compliant.

The WEGS in the initial release of the Core Grid Middleware will accept workflows written in Web Services Flow Language (WSFL) and Scufl. However, NEReSC is already evaluating the possibility of providing support for the Business Process Execution Language for Web Services (BPEL4WS) in addition to WSFL.

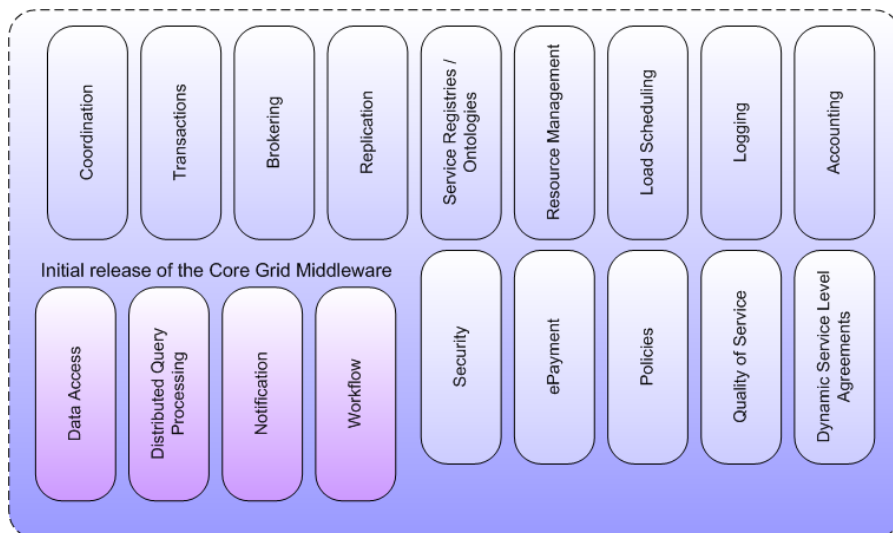


Figure 2. A potential list of OGSA services

NEReSC is planning to follow the work of the GGF working groups in the important area of workflow/business process composition and provide support for the Grid workflow standard when that is made available.

Notification

The expected dynamic and distributed nature of Grid Applications means that a facility is necessary for informing interested parties of changes in data, Grid Service status, application-specific events, etc. The Notification Grid Service (NGS) is based on the one produced within the myGrid project. The myGrid Notification Service provides a way for myGrid Services to publish events and/or register interest in published events.

Like the workflow enactment engine, the myGrid Notification Service is based on Web Services and so it is being adapted to be OGSi compliant. The OGSi specification describes a notification portType and Service Data Elements (SDEs) through which NGS will provide the myGrid Notification Service functionality.

Future Work

The work being carried out to port the myGrid Workflow Enactment and Notification Web Services to be Grid Services will develop valu-

able knowledge about the issues and effort required to achieve this conversion—something that will be useful to the many UK e-Science projects planning to carry out this transition at some point.

After the initial release, new services will be added to the core set over time, in response to requirements and availability. The main candidates will be the services specified by the OGSA standardisation activity (Figure 2), but it may be necessary to pre-empt standardisation in order to meet project requirements. In most cases “best-of-breed” services will be identified and integrated, provided that an open-source implementation is available. However it may sometimes be necessary to develop services if there is a requirement and no implementation exists.



Prof. Paul Watson



Dr. Savas Parastatidis