Building the e-Science Grid in the UK: Grid Information Services

Rob Allan, Dharmesh Chohan, Xiao Dong Wang, Andy Richards,
CCLRC e-Science Centre and Grid Support Centre, Daresbury Laboratory
Mark McKeown, Grid Support Centre, University of Manchester
John Colgrave, IBM Hursley Laboratory
Matthew Dovey, Oxford e-Science Centre
Mark Baker, Distributed Systems Group, University of Portsmouth
Steve Fisher, Rutherford Appleton Laboratory

Abstract

In this paper we review work that is being undertaken to evaluate Grid Information Service technologies and attempt to define what services and related schemas are required for a production Grids. Much of the paper is based on a report from the GIS2003 workshop held at NeSC 24-25/4/2003 [1]. We identify a need to support both static and dynamic information, and discuss what is happening in this area with UK Grids. Finally, we discuss additions to the existing schemas that would be required to support users and applications, in addition to resources.

1. INTRODUCTION

The Grid provides an infrastructure that can provide seamless, dependable, consistent, and pervasive access to distributed resources, services, groups or virtual organisations. All these components that make up the Grid can be sources and/or sinks of information, which in turn can be dynamic and/or static in nature. Key to using the Grid, or any other wide-area infrastructure is the utilisation of this information in a range of increasingly complex ways, for example from advertising simple capabilities service, to annotating and storing data that allows sophisticated workflows to be developed.

A key component of any distributed system is its so-called Information, or Directory, Service. This service typically holds information that has been advertised by an available service with that environment, and can be interrogated or searched through by a client wanting to use services. Often also the Information Service will hold addition information that can by the client to directly interact, or bind, to service that it wants to use.

Typically the interaction with Information Services involves queries via "standard" APIs for publishing and searching, and working with some definition of data stored on the server.

Each distributed system has its own Information Service, which in turn has different APIs and data definitions. In addition, the service may provide static and/or dynamic information that can be provided synchronously or asynchronously. These factors, plus others such as holding customised information and use of different versions of the same Information Service makes working with one or more Information Services from a single or cross environment perspective problematic, if not impossible and thus makes the investigation of this area a key step for the successful deployment and use of infrastructures such as the Grid.

This paper is broken up into sections. In sections 2 we discuss and detail the status and current Information Services. In section 3 we look at data and data model that are used with Information Services. In section 4 we explain what is happening within the e-Science support centre with regards to Information Services and the associated data. Section 5 outlines potential future directions for information Services, and in Section 6 we conclude.

2. CURRENT GRID INFORMATION SERVICES

The discovery and use of services and service data, is a difficult and complex task, especially in large scale, open distributed systems such as the Grid, due to:

- The global extent and the large number and variety of services advertised,
- The diverse range of repositories used to register and publish the services and their associated data (protocols),
- The variety of methods use to describe the actual data (semantics),
- Problems associated with searching the repositories and aggregating data into knowledge that can be used for reasoning purposes (different levels require different search methods).

There are currently three main-stream technologies used:

- LDAP (MDS);
- Relational (R-GMA)
- UDDI and extensions.

2.1 LDAP

The Lightweight Directory Access Protocol (LDAP) has an object oriented data model, and is an open standard and is thus vendor independent. Directories can be distributed and replicated to provide scalability and reliability. LDAP servers can provide security through granular access control and secure authentication mechanisms. Data access to servers is via standardised Internet protocols. There is a stable open source implementation, OpenLDAP [5] that MDS uses with referral and local information providers to build a hierarchical GIS.

LDAP has shortcomings. There is an inflexible data model. So that an LDAP entry is an entry is an entry, where the information model intends to map reality. Thus, there needs to be an entry for a person and a different entry for a computer. The LDAP query language is weak. It cannot give results based on computation on two different objects in a structure. There is no join operation either. With LDAP you have to have to anticipate the questions that you want to ask the server. LDAP has no historical perspective, so time stamps need to added by another layer of middleware

2.1.1 LDAP and MDS

MDS version 1 was based on a centralised LDAP server. It was based on a push model and was fairly limited in its capabilities and features. MDS version 2 was based on a decentralised LDAP server. It provided a push model, with a soft-state protocol for GRIS/GIIS registration, and caching for performance and reliability. The MDS index service was based on GRIS and GIIS. MDS version 3 is now known as the Monitoring and Discovery System. It is designed for use with the emerging Grid Services. It can cope with a variety of data, for example from core MDS v2, GLUE Computing elements, GRAM. And others. There are higher-level tools associated with MDS3, including an Index server similar to MDS2 GIIS, command line clients, C bindings, and a GUI browser.

2.2 Relational

2.2.1 RDBMS and XML DBMS

Databases systems are optimised for fast reads and writes; data is stored a tabular form and typically has tree structure and relationships. Databases support complex queries and join operation on tables. An important feature is data resilience via features such as roll back "undo".

Products, such as Oracle Gateway [6] can be used to distribute "systems". However, typically it is not easy to distribute data across multiple vendors at present, without the use of special wrappers, or proxies. Obviously it would be possible to use RDBMS as LDAP backend server.

2.2.2 R-GMA

The European Data Grid (EDG), with its Relational Grid Monitoring Architecture (R-GMA) [7] are providing a relational view over multiple information sources. R-GMA is a relational implementation of the GMA specification from the GGF, a publish -subscribe architecture. R-GMA is concerned with both information and monitoring. R-GMA associates a time-stamp with all information, and supports a powerful data model and query language. R-GMA is not a general distributed RDBMS system, but a way to use the relational model in a distributed environment where global consistency is not important.

R-GMA is able to announce, publish and collect information. There is also a separate registry and schema, as there are different requirements for distribution and replication.

R-GMA provides different kinds of producers to cater for diverse needs. They are distinguished for example by their ability to perform streaming or perform relational join operations, or by their resilience .

2.3 UDDI

UDDI, the Universal Description Discovery and Integration is a Web Services registry standard [6], which provides a data model for describing businesses and associated services and an API with a useful but somewhat limited query language.

There is also the Universal Business Registry (UBR), which is a global public registry for advertising the available services. In addition UDDI can be set up as a private or community registry.

Originally there was a UDDI consortium; now OASIS Technical Committee [9] is determining the specification. UDDI version 3 was promised for late 2002, and version v 4 is at the discussion stage.

An e-Science project maps onto a business in a reasonably straightforward way. Assertion mechanisms can ensure information published to the UBR is accurate and it could therefore be a useful way to publish electronic information about access to e-Science services worldwide.

2.3.1 UDDI – Related Technologies

2.3.1.1 WS Inspection Language

WSIL [10] provides XML document that acts as an intermediate point to a service description, can reference UDDI or WSDL or WSIL. WSIL can be used to gather service information about a deployed service that is not published in a UDDI server, but residing in a known location, typically associated with a Web server.

2.3.1.2 ebXML

ebXML [11] is a competitor to UDDI and the OASIS TC. It is registry for e-Business objects. There are variety of mappings, for example QSL and Xquery [12]. ebXML include information life cycle management, such as submitted and approved. It also has features like policies and audit trails.

2.3.1.3 *Extended UDDI Servers*

myGrid [13] is focussed on service discovery, augmented service descriptions using metadata. Providers may adopt various ways of describing their services; access polices, or contract negotiation details. In myGrid provides a system that allows multiple users, not just publishers, to add metadata to service descriptions;

The metadata may be structured according to published ontologies, so that multiple users, especially in the case of a public registry; can interpret it unambiguously;

myGrid aims to create an information model, i.e. an ontology unifying not only UDDI and WSDL descriptions, but also general metadata attachment, as well as

DAML-S [14] and BioMOBY [15] style semantic annotations, to provide a uniform way of querying and navigating service information.

The data model relies on RDF Triples (subject, predicate, object) as the means to represent all the information in a uniform manner. This information is stored in a triple store, which can be queried uniformly through the use of a query language such as RDQL [16].

UDDIe from Cardiff University [17] is attempting to deal with "missing" UDDI services, and similarly, service documents that are out of date. In addition this project is investigating ways to extend search capabilities, such as finding "partial" matches to services. In the longer term this projects also aims to study information provenance, such as audit trails or logging capabilities.

3. DATA AND DATA MODELS

Data is increasingly being marked up in XML, which is a useful step, but there needs to be a unified way of describing the entities and elements that we are interested in. Along with ways of annotating the data, there is a need for tools to manipulate the information (translate or query), and also potentially a repository to store the information. There is a need for an independent data modelling technology.

We should be able to use the same schema, there are two emerging for use with the Grid: GLUE and CIM.

3.1 Grid Laboratory Uniform Environment (GLUE)

GLUE aims to provide interoperability between EU and US HEP Grid projects. It is a common schema to describe Grid resources:

- Compute Elements,
- Storage Elements,
- Network Elements.

There are GLUE mappings to LDAP (MDS), XML, and SQL (R-GMA).

3.2 Common Information Model (CIM)

CIM [4], from Distributed Management Task Force, is an industry approach to enable the management of "real world" objects. There are four parts to CIM:

- Modelling language and syntax,
- Management schema (core, common, extension),
- Protocol to encapsulate syntax and schema (XML/HTTP),
- Compliance document.

3.3 Tools

Open and standard tools are emerging as XML is taken up widely; including those to query and translate XML data.

The emerging query languages are Xpath [18], Xquery [19], and XQL [20]. The de facto mechanism for translating between different schemas is XSLT [21], which is based on Xpath. Also, specialised data stores are being developed, such as the RDF Triple store, which is specialised database for RDF. RDF supports RDQL, the associated query language.

4. E-SCIENCE INFORMATION SERVICE TESTBED

4.1 UK hierarchical MDS and InfoPortal

Some tools, based on the Globus MDS, have been developed in WP2 of the level 2 of the UK e-Science Grid deployment, which is managed by the Grid Engineering Task Force. Further detailed information and deliverables are at the following at:

- ETF, http://www.grid-support.ac.uk/etf
- ETF Directory Services Working Group WP2, http://esc.dl.ac.uk/MDS
- InfoPortal, http://esc.dl.ac.uk/InfoPortal

We are following steps to establish a suite of Grid Information Services to facilitate resource discovery, selection and operational monitoring. To do this we are:

1. Setting up a hierarchical resource directory service based on Globus MDS2.2;
2. Assisting users to maintain their information on line and tune the system;
3. Identify additional schema requirements (e.g. for additional job manager information).
4. Providing Web-based tools to view resource-centric structure and content based on MDS and other sources (active resource table);
5. Providing Web-based tools to view site-centric content based on MDS and other sources (active UK site map);
6. Integrating with Network Monitoring service [22];
7. Integrate with Grid Integration Testing service [23];
8. Provide programmatic Web service interfaces to above information.

The UK e-Science MDS has a top-level (Tier 0) node maintained by the Grid Support Centre at ginfo.grid-support.ac.uk with VO name "UK e-Science". We are using standard port 2135 for all LDAP traffic. There are currently 4 tiers in the MDS hierarchy. In order for this to function correctly it was essential to synchronise system clocks by using NTP servers in all tiers.

In addition to the Globus command line, we are developing a number of Web Services as part of InfoPortal. Currently users can invoke an MDS query to download a complete snapshot of the status of an individual resource via a Web Service by sending its fully qualified domain name. Alternatively a cached copy of the same data can be obtained (the InfoPortal caches snapshots once per hour). In addition, services are provided to return the static XML data about individual resources. Future work will address more complex queries as a way to support brokering and scheduling.

4.2 Extending Schemas for e-Science

In a separate paper [24] we proposed XML schemas for Grid users, applications and resource descriptions, which can be used for authorisation and contact purposes, discovery and for mapping applications to resources. Such schemas could enable a range of services including Google-style searches for Grid applications across scientific disciplines with appropriate information on application functionality, input-output, QoS requirements, licensing, user authorisation and install base. The schemas are being developed in the CCLRC IeSE project [25], for instance Web Services are provided in InfoPortal to return static data about individual resources encoded in XML. These schemas effectively could be used to extend the UDDI schema for e-Science projects with the use of unique keys. We have suggested this as a mechanism to provide a semi-static information base complementary to the Globus MDS or other dynamic source. They are a first step in defining full Grid facility metadata.

5. FUTURE DIRECTIONS

Tools and standards appropriate for Grid Information Services can be classified as follows:
Emerging:

- *Data models – standards are appearing e.g. GLUE/ CIM but need to be agreed;*
- *Data providers and repositories – using the likes of MDS/R-GMA/service-data/UDDI/ RDBMS/ XML DBMS.*
- *Rapid changes are being made to take up of functionality that is missing.;*
- *Semantic markup is being introduced in a limited way.*

Being Studied:

- *Producer-consumer interactions – synchronous/asynchronous, static vs. streams;*

- *Performance and stability – being explored, e.g. for LDAP;*
- *Security – typically skipped, but role-based authorisation models will be important for an extended GIS including people, applications and data;*
- *QoS/SLA – starting to be studied;*
- *Scaling – still to be investigated.*

Future:

- *GIS as the basis of user management, brokering and scheduling tools;*
- *Federation of services – not discussed or handled yet;*
- *Contract negotiation based on service description and QoS – not even discussed yet.*

6. CONCLUSION

We believe that given the number of active projects in the UK developing Grid Information Services or using them, we should be in a good position to address some of the above topics and make better use of the available tools and schemas in a coordinated way. A full Grid information system is likely to require a combination of streamed data sources and semi-static information, some of which was described above. It should however have a relatively simple API or small number of APIs enabling tools to be developed to support complex queries and provide meaningful information to brokers and other application-level Grid information consumers.

Acknowledgements

We acknowledge all speakers and attendees of the GIS2003 workshop in Edinburgh 24-25/4/2003 and staff of the UK National e-Science Centre (NeSC) who helped with the organisation.

References

[1] R.J. Allan, M. Baker, R. Boakes & S.M. Fisher, *Grid Information Systems 2003* Technical report following GIS2003 workshop (Edinburgh 14-15/4/2003)

[2] MDS, http://www.globus.org

[3] GLUE, http://www.cnaf.infn.it/~sergio/datatag/glue/

[4] CIM, http://www.dmtf.org/standards/standard_cim.php

[5] OpenLDAP, http://www.openldap.org/

[6] Oracle Gateway, http://www.rainingdata.com/products/connect/d3oraclegateway/

[7] Cooke et al., *Relational Grid Monitoring Architecture (R-GMA)* – this conference and http://www.r-gma.org/

[8] UDDI, http://www.uddi.org

[9] OASIS, http://www.oasis-open.org

[10] WSIL, http://xml.coverpages.org/wsil.html

[11] ebXML, http://www.ebxml.org/

[12] Xquery, http://www.w3.org/XML/Query

[13] myGrid, http://www.mygrid.org.uk

[14] DAMS-L, http://www.daml.org/services/

[15] BioMOBY, http://www.biomoby.org/

[16] RDQL, http://www.hpl.hp.com/semweb/rdql-grammar.html

[17] UDDIe, http://www.cs.cf.ac.uk/User/A.Shaikhali/uddie/

[18] xPath, http://www.w3.org/TR/xpath

[19] xQuery, http://www.w3.org/XML/Query

[20] XQL, http://www.w3.org/TR/2003/WD-xquery-20030502/

[21] XLST, http://www.w3.org/TR/xslt

[22] M. Leese, Network Monitoring, Proc. AHM 2003 (Nottingham 2-4/9/2003)

[23] D.J. Baker, S. Cox, J. Hillier, M. Mckeown & R.F. Fowler Building the e-Science Grid in the UK: Providing a Software Toolkit to Enable Operational Monitoring and Grid Integration, Proc. AHM 2003 (Nottingham 2-4/9/2003)

[24] R.J. Allan, X.D. Wang, A.J. Richards, D. Chohan & M. McKeown *XML* Schemas

for e-Science Projects, Grid Users, Applications and Resources http://esc.dl.ac.uk/Papers/Schema/schema.pdf
[25] IeSE: the CCLRC Integrated e-Science Environment, http://esc.dl.ac.uk/IeSE/about.html