THE NERC DATAGRID: "GOOGLING" SECURE DATA

Bryan Lawrence[2], Ray Cramer[3], Marta Gutierrez[2], Kerstin Kleese van Dam[1], Siva Kondapalli[3], Susan Latham[2], Roy Lowry[3], Kevin O'Neill[1], Andrew Woolf[1]

[1]CCLRC e-Science Centre
[2]British Atmospheric Data Centre
[3]British Oceanographic Data Centre

Abstract

The NERC DataGrid project began in September 2002. The key aim of the project is to provide access to data holdings that lie in a (possibly very) loosely coupled federation of sites sharing at the minimum a set of common discovery, authentication, authorisation, and access protocols. While it is anticipated that much of the data will be publicly available, the NDG aims to provide secure access to protected data where trust relationships exist between data providers. The NDG design does not assume one central repository for all metadata, nor does it assume common policies (although it does require a common policy framework). However, following the "search engine" philosophy, search engines will be free to harvest publicly available "discovery metadata" using standard digital library protocols to produce high performance discovery engines. The NDG has a working prototype discovery engine, and a software infrastructure that is designed to allow migration of services to new grid architectures as easily as possible.

1. INTRODUCTION

The advent of whole earth system science has meant that the environmental science community needs better tools to discover, share and utilise data with maximum format transparency. The NERC DataGrid (NDG) is being designed and built in response to that need within the UK academic environmental science community, but the technology and design is intended to be completely scalable to the wider environmental community.

In the first instance, NDG development is being driven by the requirements of the atmospheric and oceanographic communities to handle both very high volume datasets (e.g. supercomputer simulations and space-based remote sensing) and very complex datasets (e.g. marine observations collected aboard cruises and/or atmospheric observations collected aboard flights). Such data are current stored (and possibly archived) in diverse locations, with a plethora of access control policies and various states of metadata. NDG is aimed at providing an infrastructure to handle interdisciplinary discovery and data utilisation from a user perspective as well as a toolbox to aid data providers to expose their data to the NDG.

In this paper we discuss the design principles in general, revise the NDG metadata taxonomy introduced in previous work, outline the architecture being implemented in phase one of the NDG, discuss the current prototype, and briefly indicate the development path planned into the near future.

2. DESIGN PRINCIPLES

The formal architecture of the NDG is discussed in Woolf et al. (2004), but among the key roles are Data Providers and Users. NDG has two main design principles associated with these roles: there should be no limit to scalability in number of users and number of data providers, and the overheads on being involved should be as minimal as possible.

Ideally, from a Data Provider point of view, NDG participation should be as easy as publishing on the web: in that case one creates a number of web pages, and one runs a web server (or uses someone elses). In the case of the NDG, given some data, a data provider needs to create NDG metadata, and run an NDG DataProvider interface (or use someone elses). From a User point of view, NDG participation

should be possible at three levels: via a web-client, via a command line, and eventually via a programmatic API.

Phase One of NDG, which should be fully deployed in 2004, does not reach all of these goals, but is aimed at deploying a system that scientists can begin to use now, with the confidence that these other aims will be achieved if the Data Providers conform to NDG metadata usage, and development proceeds along the planned trajectory.

3. METADATA TAXONOMY

In previous work (Lawrence et. al., 2003, O'Neill et. al., 2003) we introduced the taxonomy of metadata that forms the heart of NDG development. Over the past year, we have further refined our understanding of the relationships between these metadata types, and this is outlined in figure 1.

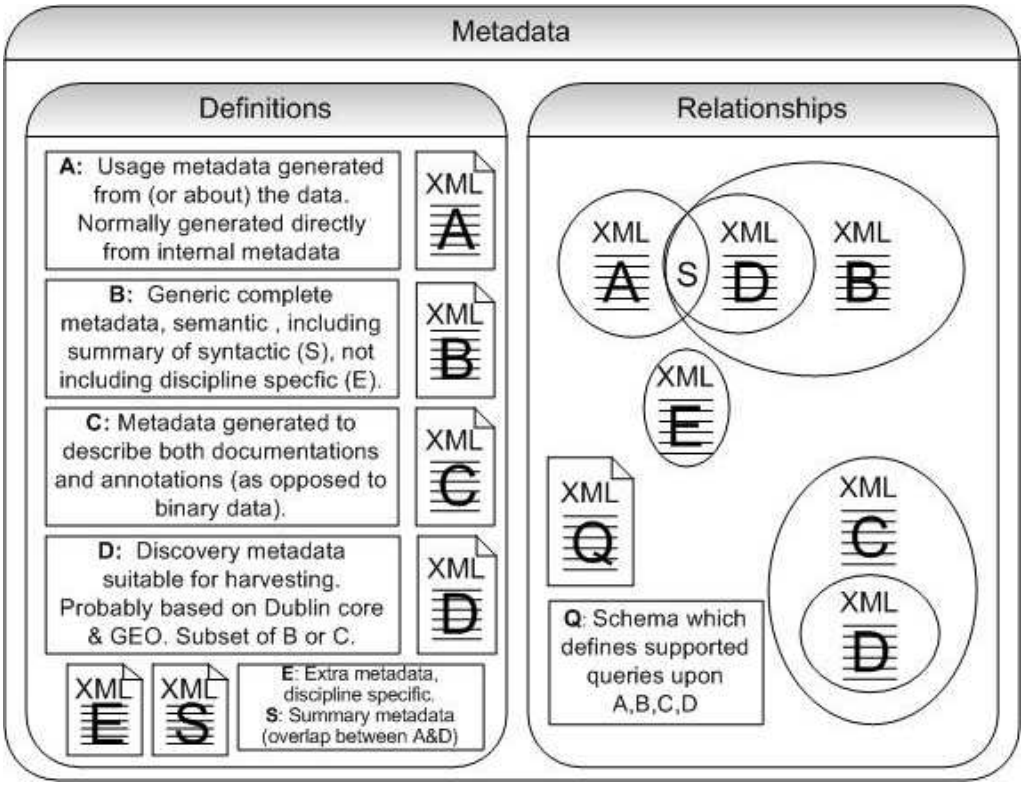

Figure 1: Metadata Taxonomy (V2004-1)

There are two further overarching design principles that applies to the NDG: Data Providers should be able to control access to their datasets with whatever granularity they desire and, where possible, standards based metadata and access protocols should be used. We believe it is the seamless transition from discovery to data access including secure access control via open standardised protocols that will differentiate the NDG from other distributed data access projects.

In phase one, the key metadata elements are publicly available discovery metadata (D), secure browse metadata (B) (which includes access control or "security" metadata), and a set of semantic "archive" metadata (A) which hides the physical data format details but exposes the geospatial and parameter characteristics of the data in a standards compliant manner (where possible, the NDG metadata development is conforming with ISO standards). The relationship between these metadata

elements is more fully explored in (O'Neill et al, 2004).

As described in more detail below, the A metadata coupled with NDG delivery services provides direct access to the data – it will not be necessary for the user to know the format of the data, and the user will be able to define an output format (from a selection of those supported by the service).

4. OVERALL DESIGN

The main components of the NDG architecture have been developed using the RM-ODP process (see Woolf, et.al, 2004). In this section we provide a brief overview of how the various components will work.

4.1 Access Control and Authentication

The NDG architecture is designed to support a federation of data providers (DPs), each of which will consist of one or more data host systems. The main characteristic of a data host system is that it hosts an archive of A and B metadata for datasets which the Provider wishes to expose. Normally a DP will also run a Gatekeeper service to constrain access to both the data and browse (B) metadata. The Gatekeeper, also known as an "Attribute Authority", consists of a set of interfaces to four key databases: a private set of recognised local credentials and their definitions, a set of publicly available credential descriptions (which ought to be a subset of the private credentials), a user database which maps local users onto local credentials, and a "credential-map" which lists a set of remote trusted gatekeepers and how their publicly available credentials map onto the local credentials.

Pseudo descriptions of how these attributes will be described can be seen in the following fragments . Access control in a B metadata record held at the BADC might read:

```
<authority="badc.nerc.ac.uk">
<accept>nerc</accept>
<deny>commercial</deny >
</authority>
```

while the credential map might include something like :

```
<trust>
<authority ="bodc.nerc.ac.uk">
<map> nerc staff </map>
</authority>
</trust>
```

The key concept is then that data access at the BADC from users who can present an attribute certificate signed by the BODC which asserts the user has "staff" access will be granted privileges commensurate with the BADC status of "nerc". It is expected that the credential map will not be automatically generated – it will be built on a sequence of bilateral trust relationships established by humans.

The reliance on bilateral trust relationships rather than an overall set of NDG roles and policies is a consequence of the NDG requirement for scalability. It became apparent early on in discussions about NDG access control that no useful global (NDG-wide) policies could be constructed that provided useful access control discrimination. Of course, it is perfectly possible that various confederations can share common policies and attribute authorities, thus allowing large domains of commonality.

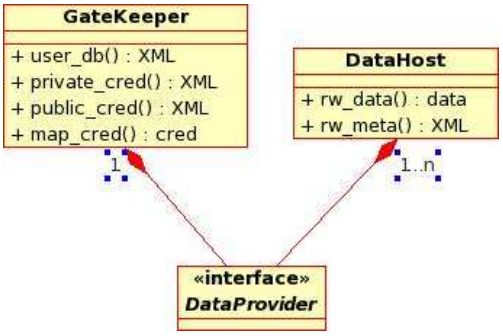

Figure 2: Data Provider Interface

A Data Provider need not itself run a gatekeeper, but it must have access to one. Thus it will be possible for a Data Provider to consist only of data hosts, utilising a remote gatekeeper, provided they trust the remote gatekeepers to maintain user and access control information. This latter configuration would be suitable in a more

tightly coupled virtual organisation. Example remote Gatekeepers envisaged could include departmental servers or shibboleth systems (Shibboleth: http://shibboleth.internet2.edu).

The Data Provider interface (Figure 2) primarily utilises the Gatekeeper in conjunction with the DataHost rw_meta interface. The latter consists of three further interfaces: an "anonymous" access to discovery metadata (D), via the digital library Open Archives Initiative (OAI) protocol (see www.openarchives.org), and two interfaces mediated by the gatekeeper; an interface to the B metadata and an interface to the A metadata (and hence directly to the data).

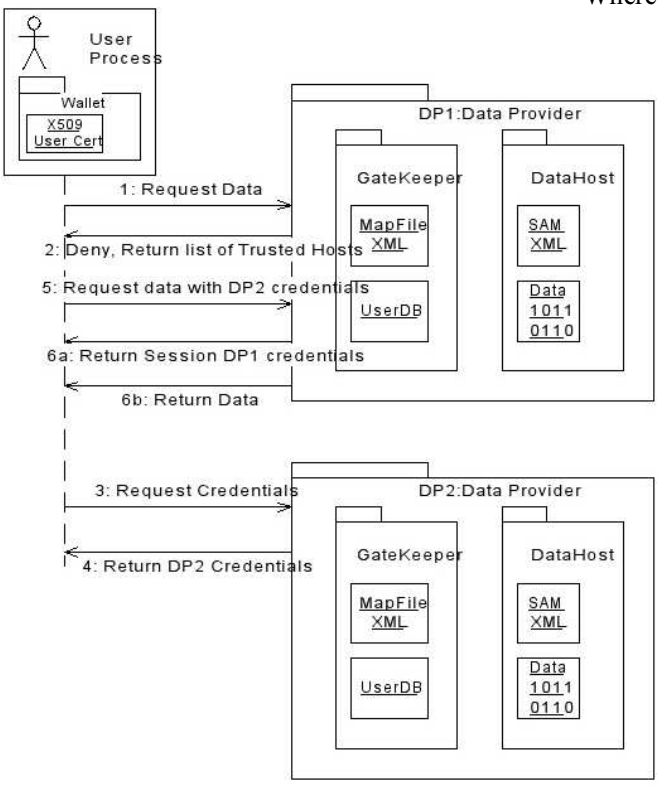

Figure 3: Access and Authorisation: Population of the User Credential Wallet

The methods by which these access control mechanisms are used are depicted in the sequence diagram (Figure 3). A key concept here is that the user software agent (user process) possesses a credential wallet, which keeps authentication and authorisation certificate, and that the user process presents certificates to NDG services when requested.

4.2 NDG Services

To orchestrate user access to data there are three services: a discovery service (essentially a search engine which consists of a user interface and an OAI repository); a browse service, which allows users to tunnel into multiple datasets at multiple locations to identify subsets of interest, and a delivery service, which optionally aggregates data and delivers it to a user process (possibly via another processing service – in phase one of NDG, a simple visualisation service will be provided). Where the user requests data, the delivery format will be one of a number of different formats supported by the semantic data model which underpins the A metadata (see Woolf et al, 2003). Note that the physical location of files could well be handled by another level of logical abstraction (e.g. the San Diego Storage Resource Broker, SRB). This is likely to be the case for at least one of the NDG Data Providers – the British Atmospheric Data Centre, BADC.

Although NDG will be building a discovery service, the design philosophy will be to allow anyone to build their own OAI based discovery service, and indeed anyone to build a browse service, provided it can provide the appropriate security credentials. to the gatekeepers.

From a user perspective the complete sequence of events from discovery to utilisation of data is depicted in Figure 4. A user begins by selecting datasets of interest from the discovery records (D metadata) harvested by the discovery service. (This is analogous to finding web pages of interest using Google). The user may then be authenticated and if access is granted, then can interact with the browse metadata services to understand more details of the

datasets of interest (utilising the B metadata). Finally, the user interacts with the A metadata to fully define the data of interest to them, and a dataset is produced and made available via the delivery mechanisms of choice.

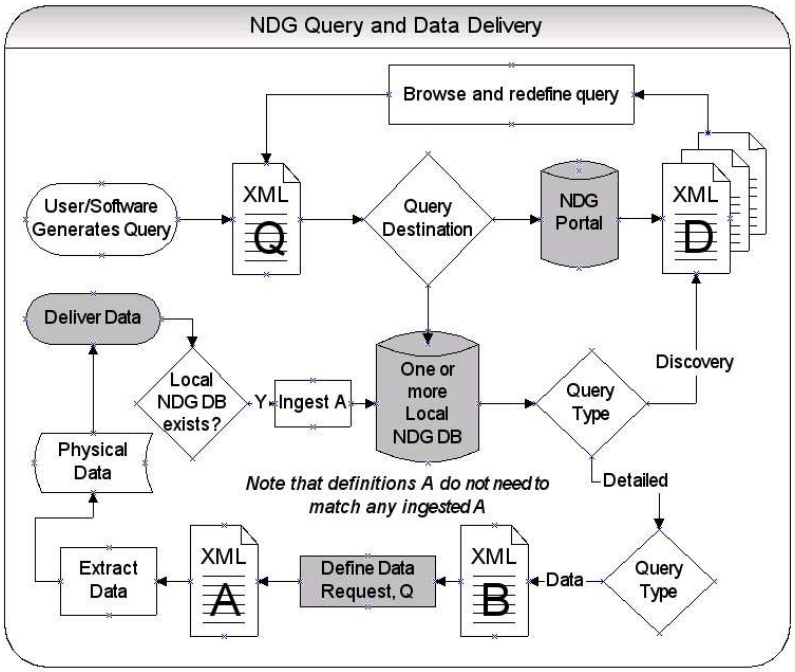

Figure 4: NDG Query Sequence. Individual discrete services are shaded.

An important characteristic of the sequence of events is that the user query is passed as an argument from the discovery to the browse service, thus ensuring the seamless transition from discovery to data utilisation. All the data arguments will be encoded as XML documents (although the databases will expose XML, they do not have to be XML-native databases).

The delivery of data to the user will need to be done in such a way that user can exploit existing software tools. To that end, we will be investigating the suitability of using the Earth System Grid (www.esg.org) OPeNDAPg protocols for secure data delivery of new datatypes directly into applications such as the Climate Data Analysis Tools (see esg.llnl.gov/cdat/).

If the user themselves runs or can ingest data to a Data Provider service (and it is hoped that large-scale users will be able to do so), the final step may be to register the data that they have extracted so that other users may be able to utilise that dataset rather than the source dataset (as may be desirable should the volumes of data be large, and significant network bandwidth issues be involved in the data transfer).

5. CURRENT STATUS

The NDG phase one implementation will be nearing completion in September 2004. The key elements of the prototype exist at the time of writing (June 2004), the major outstanding issue is an implementation of the access-control model. Meanwhile, access control simply utilises local user databases.

NDG phase one is utilising the NASA Global Change Master Directory's Directory Interchange Format (DIF) for discovery. Future versions of NDG will use ISO19115 discovery metadata.

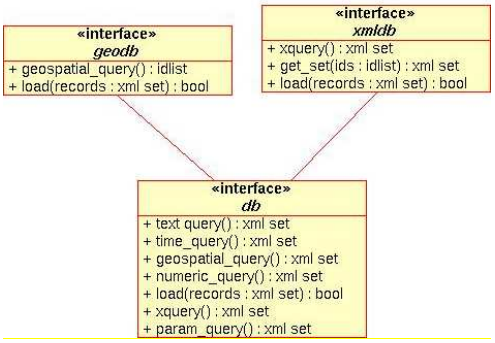

Figure 5: Database interface exposing operations which are served by two underlying DB implementations (following ideas from Liu et.al, 2003).

Phase one is implemented using open source databases. While in the long run, individual Data Providers may implement commercial databases for performance reasons, the NDG philosophy is that NDG Data Providers should not have to pay to take part – this means we cannot rely on commercial software solutions.

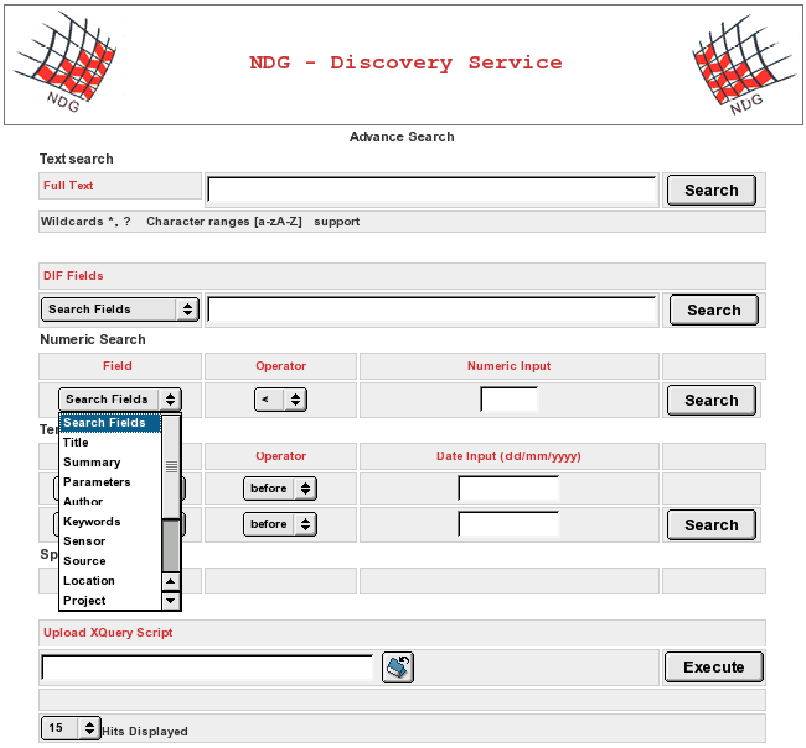

Figure 6: Advanced Search Interface. Note the drop down list displaying the option to search on various DIF parameters and the support for direct xquery expressions.

NDG phase one does not allow the user to specify geographical searching at discovery. This is because the metadata is currently stored in an XML-native database (exist, see exist.sourceforge.net), and geographical indexing is not yet available. If geographical indexing is not available soon in an open source XML database, we intend to replace our database interface with an interface that exposes both a relational database holding the geographical indexes and the other data in the current XML database (see Figure 5). When this is done, our GUI for discovery will be able to support efficient geographical queries as well as an xquery interface.

NDG provides two starting points for discovery. The simple search interface to NDG data simply exposes a text box search, similar to that of Google. The user has a number of configuration options, however, including the layout (order by title or data centre) and the amount of information returned (titles only or titles and abstracts). Future versions will allow the user to select a ranking algorithm via the advanced search page, meanwhile the advanced search page simply exposes the structure of the underlying DIF record for parameter based searching (figure 6).

When the user has selected a dataset, they can choose a number of options: they can explore the full discovery (D) metadata record, choose to examine the browse (B) metadata, or go to a data selector based on the A metadata (Figure 7) and choose either to obtain data or a plot (produced using CDAT).

We have based the development of our data selector on the excellent Live Access Server (http://ferret.pmel.noaa.gov/Ferret/LAS/ferret_LAS.html, Hankin et al, 1998). We have not used the LAS itself for technical reasons, but the LAS concepts and interfaces have and will continue to inform our development as we believe it to be the reference implementation of best practice in this area.

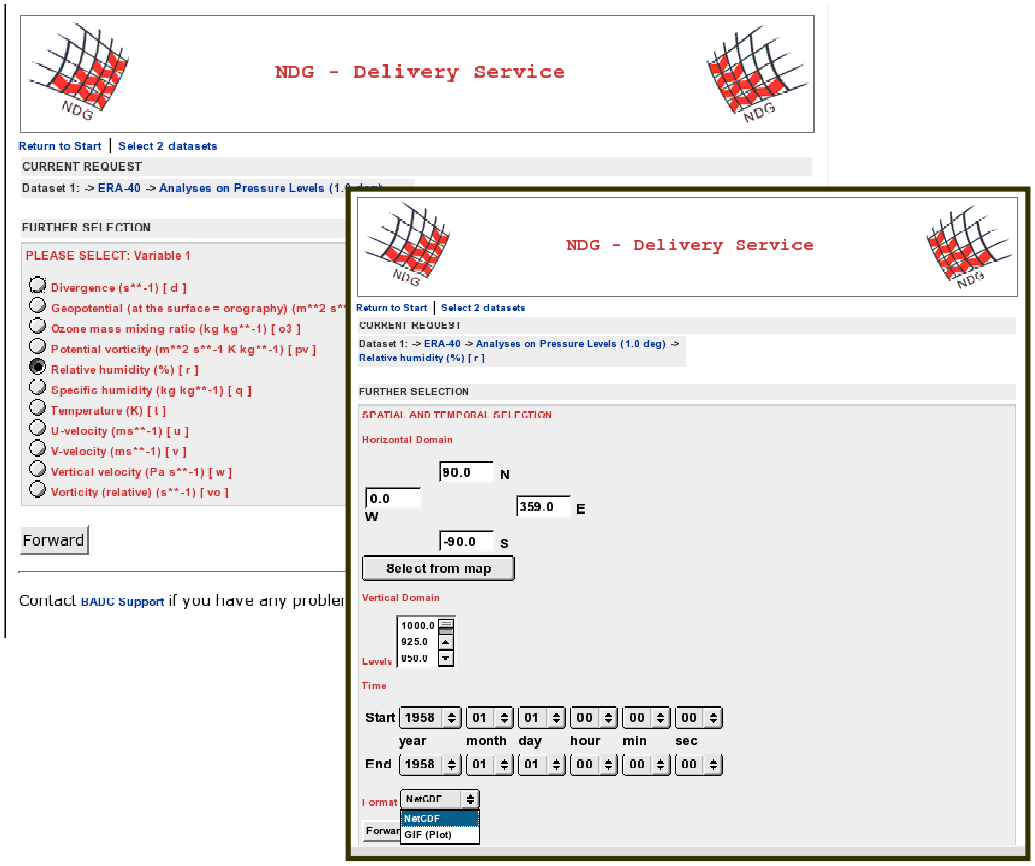

Figure 7: Variable, spatial and temporal selection in the data extractor (two web pages overlaid). Note the option to return either netcdf or a GIF plot of the data. The ability to configure plot and other options will be added shortly.

The data selector web service has already been deployed standalone at the BADC, and is being used to interface to many TB of simulation data. Major underpinning developments for this have included new code for direct support for the PP format that the Met Office uses for climate and numerical weather prediction, and indirect support via caching for atmospheric model data stored in GRIB format spectral coefficients. In this latter case, the user is presented with an interface that exposes the data as if it were in a standard latitude/longitude/height grid, and only converts the data from the spectral representation when the data is needed. This means we can store 4TB online and present it as a virtual 18TB dataset. Work on configuring the workflow at the backend of this service to exploit grid computing is underway.

6. SUMMARY AND FUTURE PLANS

Early on in the development of the NDG, it became obvious that to exploit the possibilities of semantic interoperation that grid technologies provide; we would need new metadata structures. At the time of writing we believe we have put in place the foundations of those metadata structures. In doing so, we have concentrated on a standards based approach, which we believe has two significant benefits: firstly, we can exploit existing best practice, and secondly, we expect interoperability will be significantly easier.

The major aims in the next period are to complete the re-engineering of our prototype to fully conform to a web-service architecture, and to implement the security mechanism outlined here.

To implement our security infrastructure, we will begin by extending the security mechanism built into the CCLRC Data Portal (Manandhar et.al, 2003), but we will be evaluating other grid security paradigms.

We then expect to start work on engineering to a complete architectural specification which we have developed based on the RM-ODP methodology (see Woolf et al, 2004). While the work done so far has not fully exploited grid technologies (in particular stateful services) the full architecture will require using such services.

7. REFERENCES

Hankin S., J. Davison, J. Callahan, D. E. Harrison, K. O'Brien, 1998: A configurable web server for gridded data: a framework for collaboration. In 14th International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography, and Hydrology, AMS, 417-418.

Lawrence, B.N., R. Cramer, M. Gutierrez, K. Kleese van Dam, S. Kondapalli, S. Latham, R. Lowry, K. O'Neill, and A. Woolf. 2003: The NERC DataGrid Prototype. Proceedings of the U.K. e-science All Hands Meeting, 2003

Liu, Z., E-P Lim, W-K Ng, and D.H. Goh, 2003. On querying geospatial and georeferenced metadata resources in g-portal. Proceedings of the third ACM/IEEE-CS joint conference on Digital libraries, 245-255.

Manandhar, A., G. Drinkwater, R. Tyer, K. Kleese. GRID Authorization Framework for CCLRC Data Portal. Proceedings of the All Hands Meeting 2003.

O'Neill, K., R. Cramer, M. Gutierrez, K. Kleese van Dam, S. Kondapalli, S. Latham, B.N. Lawrence R. Lowry, and A. Woolf. 2003: NDG Metadata: interfaces between discovery, browse and semantic utility. Proceedings of the U.K. e-science All Hands Meeting, 2003.

O'Neill, K., R. Cramer, M. Gutierrez, K. Kleese van Dam, S. Kondapalli, S. Latham, B.N. Lawrence R. Lowry, and A. Woolf. 2004: A specialised metadata approach to discovery and use of data in the NERC DataGrid. Proceedings of the U.K. e-science All Hands Meeting, 2004.

Woolf, A., R. Cramer, M. Gutierrez, K. Kleese van Dam, S. Kondapalli, S. Latham, B.N. Lawrence, R. Lowry and K. O'Neill. 2003: Data Virtualisation in the NERC DataGrid. Proceedings of the U.K. e-science All Hands Meeting, 2003.

Woolf, A., R. Cramer, M. Gutierrez, K. Kleese van Dam, S. Kondapalli, S. Latham, B.N. Lawrence, R. Lowry and K. O'Neill. 2004: Enterprise Specification of the NERC DataGrid. Proceedings of the U.K. e-sicence All Hands Meeting 2004.