University of Glasgow | Experimental Particle Physics

GLAS-PPE/2009-23
19th November 2009

Department of Physics and Astronomy
Experimental Particle Physics Group
Kelvin Building, University of Glasgow
Glasgow, G12 8QQ, Scotland
Telephone: +44 (0)141 330 2000 Fax: +44 (0)141 330 5881

Spatial Data e-Infrastructure

C. Higgins[1], M. Koutroumpas[1], R.O. Sinnott[2], J. Watt[2], T. Doherty[2], A.C. Hume[3], A.G. Turner[4], D. Rawnsley[5],

[1] EDINA National Data Centre, University of Edinburgh, Edinburgh, EH9 1PR

[2] National e-Science Centre, University of Glasgow, Glasgow G12 8QQ ..

[3] Edinburgh Parallel Computing Centre, University of Edinburgh, Edinburgh, EH9 3JZ

[4] National Centre for e-Social Science, University of Leeds, Leeds, LS2 9JT ..

[5] MIMAS, University of Manchester, Manchester, M13 9PL

Email: r.sinnott@nesc.gla.ac.uk or t.doherty@physics.gla.ac.uk

Abstract

This paper examines overlap between e-Science and geospatial communities using work undertaken as part of the JISC funded Secure Access to Geospatial Services (SEE-GEO) project. Working from a position that open standards provide the best means of achieving interoperability, the case studies demonstrate the use of Grid technology and open geospatial standard interfaces to realise classic spatial data infrastructure scenarios. These examples are used to illustrate just how e-infrastructures enable seamless, security driven spatial data access at the national, regional and global scale.

Introduction

There is an increasingly vast amount of data that can be or is geographically referenced and available via the internet. Providing the computational infrastructure, the technologies, the tools, the resources necessary to make the most of this data has been the subject of significant amounts of effort. This effort is genuinely cross-disciplinary in its impact and has implications to the way we organise; in business, government and academia. In recent years, the formal framework for the development and use of such data have become known as Spatial Data Infrastructure (SDI). SDI initiatives are underway at scales ranging from the local to the global.

In association with the UK Joint Information Systems Committee (JISC) funded Secure Access to Geospatial Services (SEE-GEO) project (http://www.edina.ac.uk/projects/seesaw/seegeo), this paper reports on work which has taken place exploring the boundary between focussed e-Science and geospatial communities. Working from the position that open standards provide the best means of achieving interoperability, this project has based its development on Open Geospatial Consortium (OGC) implementation specifications. Widely used Grid middleware have been used including: the Open Grid Service Architecture Data Access and Integration (OGSA-DAI) software (http://www.ogsadai.org.uk); the Globus toolkit version 4 (GT4) (http://www.globus.org/toolkit) and the GridSphere portal framework (http://www.gridsphere.org). Additionally, a novel security solution has been implemented based on the Open Middleware Infrastructure Institute (OMII-UK) Security Portlets simplifying Access to and Management of Grid Portals (SPAM-GP) project (http://www.nesc.ac.uk/hub/projects/omii-sp).

The case study that is presented is based on linking geographically referenced data from the UK Census Programme served out by the JISC funded UK national datacentres: EDINA (http://edina.ac.uk/) and MIMAS (http://www.mimas.ac.uk/). This provides an example of how distributed data may be linked at run time using open standards on UK e-Infrastructure to produce linked data and geographic map visualisations. The tools configured enable the effective use of distributed computational resources and have a widespread, generic utility. The case study is an example of a common activity in an SDI context.

Section 2 provides a brief overview of the service oriented architectures that underpin SDI and discusses key areas of overlap with e-Science. Many of the data access, processing services and other computational resources that SDI contributors may wish to make available have access control restrictions for commercial, confidential or licensing reasons. Section 3 reports on work within SEE-GEO exploring the use of a variety of security technologies, including Shibboleth (http://shibboleth.internet2.edu/), Globus Security Infrastructure (GSI) (http://www.globus.org/security), PrivilEge and Role Management Infrastructure Standards Validation (PERMIS) project (http://www.permis.org), and Globus Virtual Organisation Membership Service (VOMS) (http://www.globus.org/grid_software/security/voms.php). These technologies may be integrated to allow e-infrastructure users access to the resources they need in a user-oriented manner, yet supporting the fine grained access control demanded by data providers. Section 4 describes these through scenarios and implementations providing secure access to geospatial services deployed on the UK National Grid Service

(http://www.ngs.ac.uk), which themselves provide access to EDINA and MIMAS resources, each of which is able to make their own local authorisation decisions.

Spatial Data Infrastructures (SDI)

One widely quoted definition of SDI is that from the SDI Cookbook (Nebert, 2004):

"the relevant base collection of technologies, policies and institutional arrangements that facilitate the availability of and access to spatial data. The SDI provides a basis for spatial data discovery, evaluation, and application for users and providers within all levels of government, the commercial sector, the non-profit sector, academia, and by citizens in general."

This definition frames SDI as a very broad concept; it encompasses more than technology and includes the social, political and economic aspects necessary to ensure that those working with spatial data are not impeded in meeting their objectives. The definition also emphasises the cross sectoral nature of SDI and it should be clear that, within the academic sector, SDI is of interest to a large number of disciplines dealing with geospatial information, including; hydrology, town planning, climatology, ecology and archaeology as well as computer science and geography.

At the global scale are initiatives such as Global Earth Observation System of Systems (GEOSS) (http://www.epa.gov/geoss/), the United Nations SDI (http://www.ungiwg.org/unsdi.htm/), and the Global SDI (http://www.gsdi.org/). At the continental and European scale, one of the most significant developments in recent years has been the passing of the Infrastructure for Spatial Information in Europe (INSPIRE) directive in 2007 (http://www.ec-gis.org/inspire/). This will make it a legal requirement for European Union member states to make available relevant and harmonised geographic information to support formulation, implementation, monitoring and evaluation of community policies. INSPIRE is the first step of a broad multi-sectoral initiative and provides the legislative basis for a European Union SDI. At the national level in the UK the publication of a Location Strategy for the United Kingdom (http://www.gipanel.org.uk/gipanel/gistrategy/) is arguably a significant step toward a UK national SDI.

Open Standards

To achieve interoperability, SDI implementation is heavily dependent upon the use of open standards. In the geospatial domain, the main Standards Defining Organisations (SDO) are the Open Geospatial Consortium (OGC) (http://www.opengeospatial.org/) and the International Standards Organisation (ISO) Technical Committee 211 (TC/211) (http://www.isotc211.org/). The OGC is an international industry consortium established in 1994. It consists of approximately 350 companies, government agencies and universities participating in a consensus process to develop publicly available interface specifications. OGC specifications support interoperable solutions that "geo-enable" the web, wireless and location-based services, and mainstream IT. ISO is a Non Government Organisation established in 1947 and based in Geneva. It is based upon national membership and has 147 participating countries. In relation to geodata standards, the most important committee is ISO TC/211Geographic Information/Geomatics which is responsible for the ISO 19000 series of standards and has as its objectives to:

- support the understanding and usage of geographic information;
- increase the availability, access, integration, and sharing of geographic information;
- enable interoperability of geospatially enabled computer systems;
- contribute to a unified approach to addressing global ecological and humanitarian problems;
- ease the establishment of geospatial infrastructures on local, regional and global level; and,
- contribute to sustainable development.

Both the OGC and ISO TC/211 work together with other relevant SDOs including the Organization for the Advancement of Structured Information Standards (OASIS) (http://www.oasis-open.org) and the World Wide Web Consortium (W3C) (http://www.w3.org/).

SDI Architecture

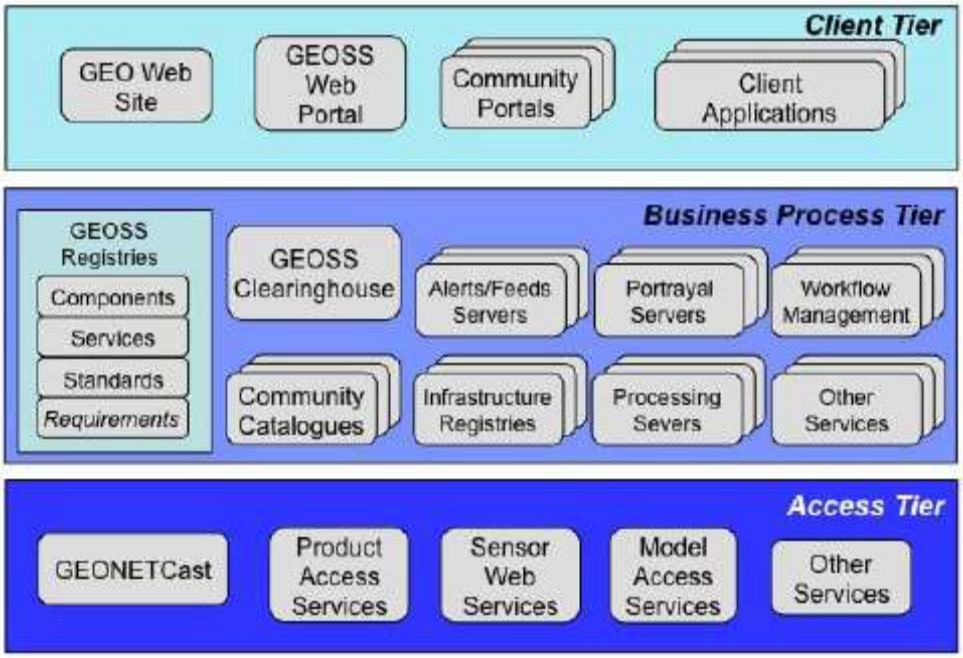

Figure 1 High Level GEOSS Architecture (Percivall, 2007)

Figure 1 is a typical SDI high level architectural overview of an SDI infrastructure. The diagram has been created in association with GEOSS (Nebert, 2008).

For simplicity and clarity, interactions between the components are not included in this diagram. Instead it focuses on showing a service-oriented architecture and makes a distinction between 3 tiers in an SDI.

1. The **client tier** contains a wide variety of both "thick" and "thin" clients capable of accessing any or all other elements using standard interfaces and encodings.

2. The **business process tier** provides distributed services that are common to the community or that perform value-added services on data. Registries are at the heart of SDI and provide the means of publishing and discovering a wide variety of different geospatial and related resources. The latter includes web services (both data access and geoprocessing), dataset metadata, schemas used for encodings, information on what standards are being used, information on who has responsibility for what, etc. Great flexibility is required of the registry technology used and careful consideration given to governance issues.

3. The **access tier** includes lower level retrieval and processing using standard interfaces for data or information from databases, sensors, or other repositories. A common requirement of SDI is access to framework datasets, e.g. roads, rivers, cadastral parcels, geology, geographic names, heights, etc.

The SEE-GEO Project

The JISC funded SEE-GEO project is a multi partner project that started towards the end of 2006 and is due to complete by Oct 2008. The two key objectives of the project are to;

- arrive at recommendations to the JISC on how to progress with the provision of secure access to distributed Geospatial Information, with particular focus on the application of Grid technology and the work of the OGC; and

- develop a number of client applications to demonstrate secure access to heterogeneous data sources hosted by the JISC funded national data centres via OGC web services integrated into standard Grid middleware.

The constituent members of the SEE-GEO consortium are: EDINA; the National e-Science Centre (NeSC) (http://www.nesc.ac.uk); the National Centre for e-Social Science (NCeSS) (http://www.ncess.ac.uk); and MIMAS.

The e-Social Science Demonstrator

Both EDINA and MIMAS are members of the Economic and Social Research Council (ESRC) Census Programme and provide services to allow users in the UK tertiary academic sector access to geospatial and census data. Through the UKBORDERS service (http://edina.ac.uk/ukborders/), EDINA holds definitive copies of the administrative and boundary framework geographic datasets including Census Output Areas (COA), the finest scale administrative output areas for the 2001 Census of huiman population. MIMAS provides access to the census aggregate statistics and postcode lookup tables through the Census Dissemination Unit (http://cdu.mimas.ac.uk/).

In 2007, the OGC initiated the Geolinking Interoperability Experiment (IE) (http://www.opengeospatial.org/projects/initiatives/geolinkie). The purpose an OGC IE is to provide the opportunity for OGC members to work together to harden candidate specifications with the intention of furthering their progress in the specification programme. The purpose of the geolinking experiment was to advance standards which separate access to framework geospatial datasets from the wide range of attributes that may be attached to them.

As SEE-GEO was conceived of as part of the JISC Grid OGC Collision programme, this IE was judged a good opportunity to engage with the OGC process and gain by working with like minded members of the international community. During 2007, under the auspices of the Geolinking IE, an e-Social Science demonstrator was created using a use case created by NCeSS (Turner, 2008).

An exemplar application was developed to demonstrate the creation of maps at run time showing the distribution of a particular statistic of interest according to the geography of interest. For example, a user may want to view long term unemployed people in Leeds according to COAs or the distribution of Gaelic speakers in Scotland according to postcode districts. Figure 2 illustrates the architecture of the implementation.

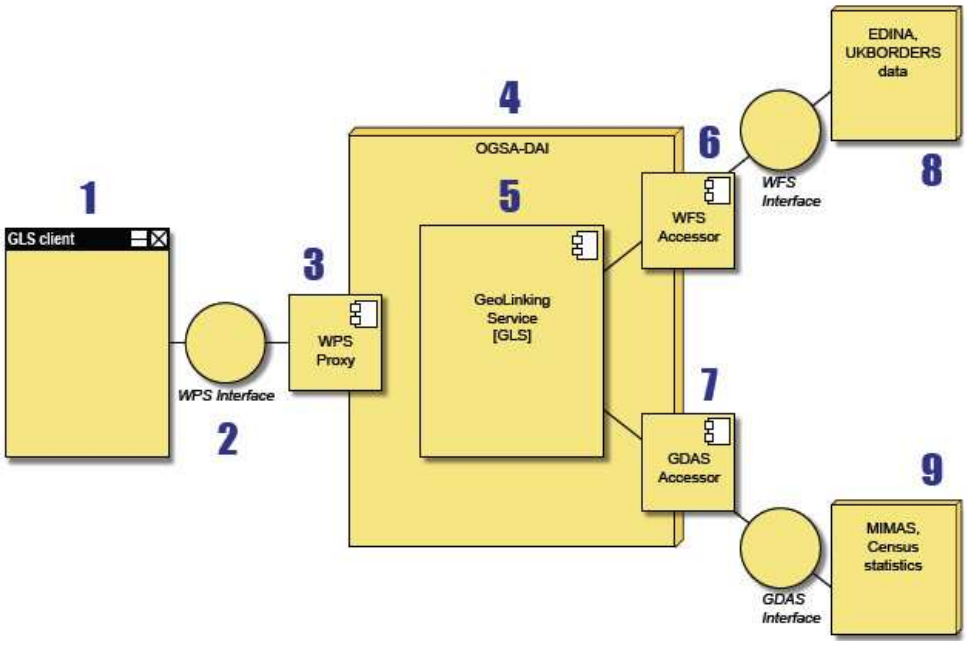

Figure 2. High Level Architecture – SEE-GEO e-Social Science Demonstrator

In this architecture, the **Geo Linking Service (GLS) client** was created by the NCeSS MoSeS (Modelling and Simulation for e-Social Science) node (http://www.geog.leeds.ac.uk/people/a.turner/projects/MoSeS/) and is used to support a wide variety of scenarios where it is useful to be able to create custom maps showing the distribution of a wide variety of health related statistics.

The OGC Web Processing Service **(WPS) interface** provides a generic mechanism to describe and web-enable a wide variety of geospatial processes. In this case, the process is linking geospatial features to attributes at run time. WPS became a full OGC interface specification in Summer 2007.

The **WPS Proxy** uses OGSA-DAI as a generic toolkit for building OGC compliant WPS. This was achieved by exploiting features of OGSA-DAI for querying, transforming and delivering data in different ways, and the features it provides for creating client applications. Since OGSA-DAI has been designed to be extensible,

users can also provide their own additional functionality. In this implementation, the **GLS** is being realised as a WPS profile. It links geographically related attribute data from a Geospatially-linked Data Access Service (**GDAS**) with geometric features from separate geospatial datasets (in this case, supplied by a Web Feature Service). We note that to support this, a common geographic identifier is a prerequisite for geolinking.

The OGC Web Feature Service (**WFS**) Specification allows the retrieval and update of geospatial data encoded in Geography Markup Language (GML) (http://www.opengeospatial.org/standards/gml). The WFS specification defines interfaces for data access and manipulation operations on geographic features, using the Hypertext Transfer Protocol (HTTP). Through these interfaces, a web user or service can combine, use and manage geodata from different sources.

The **GDAS** was an integral part of geolinking and a candidate OGC specification. GDAS delivers geographically related data (not geometries) in a simple XML format that can be used in a variety of ways. In this case, a GDAS stream is being merged with a GML stream to create an amended GML stream incorporating the additional attribute information provided by the GDAS.

The **UKBORDERS** resource is part of the ESRC Census Programme. This online service provides access to a wide variety of digitised UK boundary datasets. An OGC WFS interface has been made available to supply COA geometries for geolinking.

The **MIMAS Census Statistics** resource is part of the ESRC Census Programme. Through the Census Disseminate Unit at MIMAS a variety of census data including UK Census Aggregate Statistics is made available. A GDAS interface has been made available to provide a variety of census statistics for geolinking.

One of the main points to take from the above is, that for any of these SDI framework datasets, there is a huge variety of potential attributes that may be usefully associated with them. It is far more efficient to provide distributed access to an authoritative source of the framework dataset via standards based web services, and allow the significant numbers of institutions and users that hold these attributes to link their data at run time, than it is to have multiple local copies of the framework datasets, with databases becoming bloated by holding both the geographies of interest and the attributes of interest.

It should be apparent that there are a very large number of SDI scenarios where this technology is applicable. As the definition of SDI provided above indicates, for this to work there needs to be broad agreement on the interoperability standards being used between the providers of the web services on top of the framework datasets, the potentially much larger number of providers of attribute data, client application developers, and the range of other tool providers. Due to the large number of vested interests and level of resourcing required, such agreement typically proves difficult to reach, which is part of the reason why legislative instruments such as the INSPIRE Directive (the basis for the European SDI) are so important.

The implementation above required the integration of key OGC interfaces into OGSA-DAI which we note is also part of the Globus Toolkit and OMII-UK (http://www.omii.ac.uk/). In this instance, we are effectively using OGSA-DAI as a

toolkit for building OGC Web Processing Services (WPS). Having created the basic geospatial building blocks, e.g. the WFS Accessor, within OGSA-DAI, the potential now exists for reuse of these tools to quickly assemble new WPS to perform a range of common, and less common, SDI related tasks.

Another benefit of using OGSA-DAI is that when the feature based approach employed in the open geospatial interoperability standards stack is combined with the ability of OGSA-DAI to take advantage of multi-processor machines, performance enhancements accrue as a parallel processing pipeline approach can be used (Koutroumpas, 2008). For example, in the implementation above, the geographic features, i.e. the administrative areas, are streamed back as GML from the WFS. Each individual feature is recognised by the GLS OGSA-DAI component and, where appropriate, has its GML augmented by an additional attribute from the data returned by the GDAS. With careful design, OGSA-DAI can take advantage of multiprocessor architectures, and process each of these activities as stages in a pipeline concurrently in separate threads.

Securing Services

When establishing the Grid OGC Collision programme, the JISC recognised at the outset that secure access to services was one of the areas where the Grid community may be regarded as more advanced than the geospatial community.

Within the OGC, most effort in the security area has taken place within the context of the Geospatial (Digital) Rights Management Working Group, the Security Working Group, and the OGC Web Services (OWS) initiatives. The ratification and publication of the GeoXACML Implementation Specification (http://www.opengeospatial.org/standards/geoxacml) was a significant development. The OWS initiatives are sponsored international, collaborative, prototyping programmes which bring together large numbers of member organisations with the intention of furthering the specification programme. They focus on interoperability problems identified by the sponsors, and in recent years there has been an emphasis on security round OGC web services. This has happened in part because of the widespread recognition of the need to provide secure access to commercially valuable, licensed or confidential resources in SDIs.

Within the UK Grid community, the creation of the UK Access Management Federation (http://www.ukfederation.org.uk) and associated adoption and roll-out of Shibboleth as the primary means of authentication in the academic sector has led to various initiatives, e.g. GLASS (http://labserv.nesc.gla.ac.uk/projects/glass/), DyVOSE (http://labserv.nesc.gla.ac.uk/projects/dyvose/), VPMan (http://www.jisc.ac.uk/whatwedo/programmes/programme_einfrastructure/vpman.aspx), GridShib (http://gridshib.globus.org/), SHEBANGS (http://www.jisc.ac.uk/whatwedo/programmes/programme_middleware/proj_shebangs.aspx), ShinTau (http://sec.cs.kent.ac.uk/shintau/) and more recently SARoNGS (http://www.jisc.ac.uk/whatwedo/programmes/programme_einfrastructure/sarongs.aspx), exploring how user-oriented Single Sign On (SSO) through Shibboleth may be extended to Grid resources. Typically, Grids make use of Public Key Infrastructures (PKI); requiring users to apply for, and look after their own X.509 certificates issued by a centralised Certificate Authority. The overhead involved in doing this, the degree of technical knowhow required in their use, and the burden placed on the user for their

care, is widely recognised (Sinnott, 2005) as being a significant barrier restricting the uptake of Grids. However, the use of Shibboleth should in principle remove this barrier since end users are only required to authenticate at their home institution.

The original vision of Grids emphasised their collaborative function (Snelling, 2007):

…involve multiple organizations and individuals, security domains, protocols, discovery mechanisms, and heterogeneous hardware, collaborating to share their resources to make the most effective use of it for their combined user communities.

As a consequence, from the outset, Grids have been built with the requirement that they support potentially short/long running collaborations using distributed resources in different administrative domains. One of the issues with the roll-out of Shibboleth and its use for accessing Grid resources remains the attributes that are exchanged from Shibboleth Identity Providers and Shibboleth protected Service Providers. The UK Federation has identified a core set of attributes based around the *eduPerson* class definition, however it is clear that different Grid collaborations will want to define their own virtual organisation specific attributes and use them for their own virtual organisation-specific authorisation policies. Through projects such as SPAM-GP, it has been shown that it is possible to create portlets supporting the definition, scoping and acceptance, and subsequent use of virtual organisation specific attributes, or more precisely virtual organisations where the services and resources that are being protected are accessible within a portal based research environment.

The SEE-GEO demonstrator has exploited these results to personalise the contents of geospatial and social science portals as well as allowing providers such as EDINA and MIMAS to define and implement their own access control policies using these attributes. To realise this and support the seamless linkage of services and security infrastructures defined for example using the role based access control solution PERMIS, the services themselves, e.g. the WFS service, have been wrapped as GT4 services which allow the exploitation of Security Assertion Markup Language (SAML) call-out capabilities. Experiences of applying the SAML call-out with Globus based Grid services are described in Sinnott (2004). We note that OMII-UK now also supports this interface.

Securing the SEE-GEO e-Social Science Demonstrator

Two approaches have been explored in securing this demonstrator.

Secure the Geolinking Portal

The GLS client created by NCeSS and depicted in Figure 3 is a portlet running in the GridSphere portal framework. Building on work undertaken as part of the SPAM-GP project, authentication at this portal is achieved through logging into a recognised (trusted) Identity Provider. This trust is enforced through the SPAM-GP Scoped Attribute Management Portlet (SCAMP). This portlet filters the identity providers in the UK Federation to allow only those that are recognise as having privileges to access the portal, and further allows the recognition of appropriate roles from those providers. This scoping is achieved through filtering the attribute acceptance policy which defines the metadata for the UK federation. Furthermore, once an authenticated user has access to the portal, the roles and privileges that are released from their identity provider are used to configure and personalise the contents of the portal itself.

This manifests itself in the NCeSS client as different subsets of statistics available for mapping to different users. For example, senior, more authorised individuals get the whole range of statistics available for mapping, including possibly sensitive statistics relating to health, and less senior users, e.g. students, get presented with a range of less sensitive statistics. This is entirely configurable and can be done a virtual organisation specific manner. Figure 3 shows the results of two different test users (shibb1 and shibb2) that have authenticated at the portal and the results of the portal configuration showing the different pull down queries associated with these privileges.

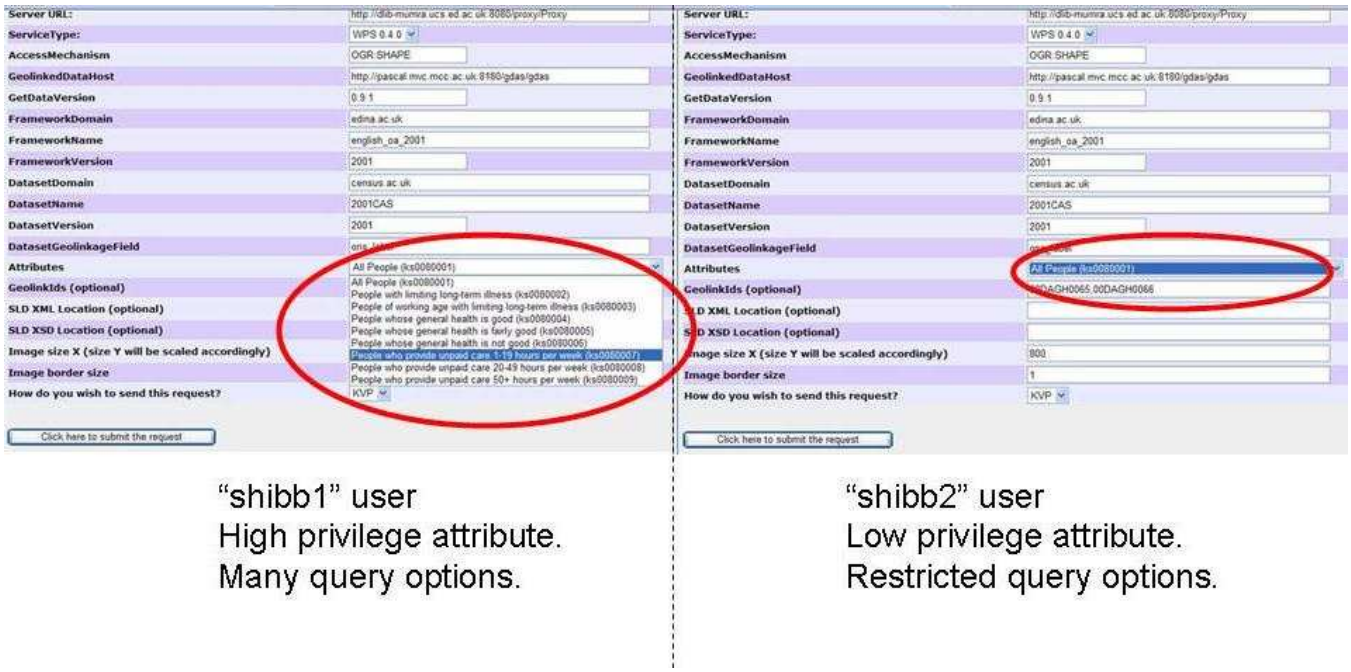

Figure 3 Privilege controlled portlet configuration

Fine grained authorisation at the Service Providers

The second approach explores the more flexible option of authorisation at the Service Providers. Thus, whilst roles released by an identity provider and accepted at the portal (service provider) can be used to provide a degree of authorisation in terms of "what you can see is what you can do", ultimately service providers such as EDINA and MIMAS will typically not simply delegate access control to a remote portal. Instead they will typically wish to define and enforce their own local policies based upon attributes that are presented to it. To support this scenario, the SPAM-GP Attribute Certificate Portlet (ACP) allows for the released roles to be signed as attribute certificates and stored in an LDAP server associated with the portal itself. Once stored, invocations or access requests to remote and protected service providers more generally, i.e. services which have defined and wish to enforce their own local access control decisions, are able, through pulling the necessary attributes from a trusted authority – namely the LDAP server associated with the portal, to subsequently decide whether a given request is a valid one or not. Figure 4 shows a schematic of the service interactions required to secure one of the backend data services (WFS) using GT4, PERMIS, MyProxy and a modified GLS client.

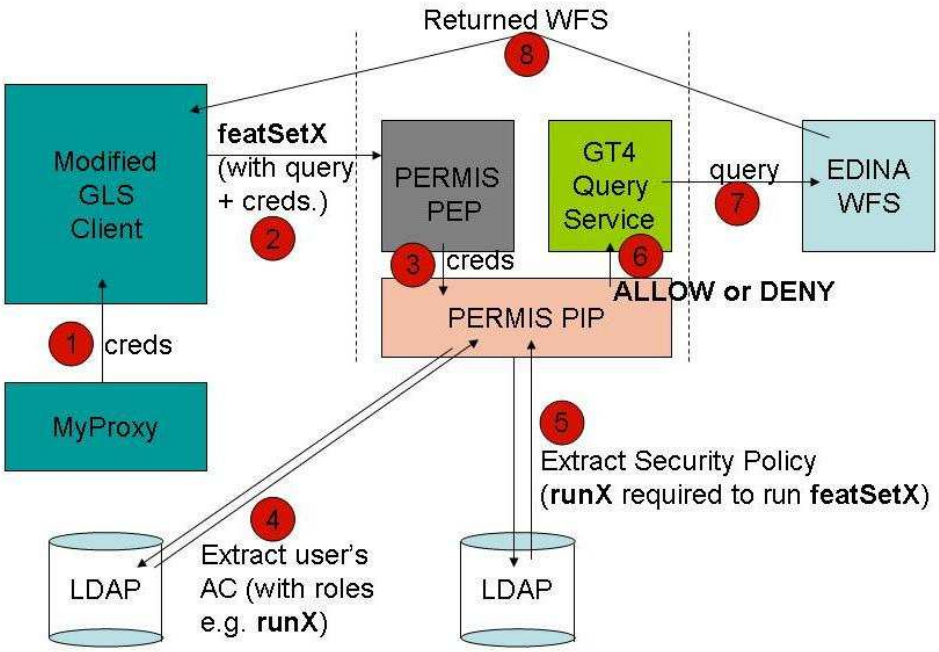

Figure 4. High level architecture – securing the Service Provider

To support this scenario, it is necessary for the client in the portal to provide information on the user making the invocation. This is achieved through exploitation of a MyProxy server at the back end of the portal. End user invocations carry with them proxy certificate information from which the remote service policy definition point can use to extract the distinguished name (DN) of the end user. Knowing the identity of the user, the portal LDAP server can subsequently be queried for the attributes associated with that particular user. This pulling of attributes is now possible through recent Open Grid Forum (OGF) authorisation standards. Other more complex scenarios including pushing or pulling of sets of attributes are also possible and will be explore in the ShinTau project in which NeSC Glasgow are also involved.

We note that a further model is also possible to define and use these attributes. Rather than using a federated model of attribute definition where the attributes are pushed and subsequently pulled from a distributed identity providers, it is also possible to agree upon a central virtual organisation specific attribute authority. One common model of achieving this is through VOMS. The pros and cons of centralised and decentralised attribute authorities for access control in a Grid environment are described in detail in Sinnott (2008).

Future Directions

Perhaps the best starting point for a discussion of the possible future direction of fruitful collaboration between the Grid and geospatial communities is to ask why the JISC funded national data centres are interested in Grid technology. Both EDINA and MIMAS exist to enhance the productivity of research, learning and teaching in the UK higher education sector. One of the ways they seek to do this is by providing core components of the UK academic SDI, including data access services on top of framework datasets, e.g. definitive administrative boundaries, topographic data, geological data. They have lots of users. In 2008, there were approximately 30000

students registered for access to Ordnance Survey GB data through the EDINA Digimap service. To fully expose a range of SDI related data access and processing web services posits a range of problems, some of which can be addressed by Grid technology and exploitation of advanced authorisation solutions.

Consider the situation where an OGC Web Feature Service on top of the most detailed topographic data in the UK, i.e. Ordnance Survey MasterMap, is made available to suitably authorised students. Firstly, there is the need for mature, reliable, interoperable solutions to the security problem as the data is commercially valuable. Secondly there is the need to be able to scale; service load varies tremendously over the academic year and purchasing hardware in anticipation of highly infrequent peak load is not an option. Furthermore, on the research side, users require the ability to collaborate by establishing virtual organisations of varying duration, and tools that facilitate dynamic resource provisioning. The latter includes the ability to be able to define chains of web services to create integrated products.

As standards defining organisations, the OGF and OGC share some common objectives. In 2007, the SEE-GEO project organised an OGF OGC Collaboration workshop at OGF20 in Manchester attended by senior representatives of both communities. In 2008, this led to a signing of a Memorandum of Understanding between the OGC and OGF. The key common objectives identified are to provide distributed data processing capabilities for geospatial data users in a way that is:

- *Transparent* - the users do not have to be aware of the exact data and computing resources they are using or the details about doing so,

- *Interoperable* - the data and computing resources can come from different sites, and

- *Scalable* - the same user model can service small computing tasks that can be done locally, as well as large computing tasks that require massive remote platforms.

Attention has initially focussed on the OGC WPS as this is the most generic of the OGC interface specifications and amenable to integration with Grid "back-end" processing environments. As well as integration with OGSA-DAI as undertaken by the SEE-GEO project, we will likely see WPS being integrated with relevant OGF specifications such as the High Performance Computing-Basic Profile (http://www.ogf.org/documents/GFD.114.pdf); the Simple API for Grid Applications (SAGA) (https://forge.gridforum.org/sf/projects/gridrpc-wg); Grid-RPC (https://forge.gridforum.org/sf/projects/gridrpc-wg) and the Web Services Resources Framework (WSRF). It is likely we will also see WPS being integrated with various workflow management tools.

References

Nebert, Douglas. D. ed. (2004): *GSDI Cookbook Version 2*, http://www.gsdi.org/docs2004/Cookbook/cookbookV2.0.pdf

Nebert, Douglas. D. (2008) *Global Earth Observation System of Systems (GEOSS): A Global SDI Framework*, http://portal.opengeospatial.org/files/?artifact_id=26314

Koutroumpas, M. and Higgins, C. (2008) *A Pipeline Processing Approach to GIS.* Proceedings from the Association Geographic Information Laboratories Europe (AGILE) Conference

Percivall, G. (2007) *Technology Viewpoint, GEOSS Architecture Initial Operating Capability*, http://www.ogcnetwork.net/IOCTechview

Sinnott, R.O. (2005) *Grid Security: Practices, Middleware and Outlook*, http://labserv.nesc.gla.ac.uk/projects/GridSecurityReport/

Sinnott, R.O. and Chadwick, D.W. (2004) *Experiences of Using the GGF SAML AuthZ Interface.* Proceedings of UK e-Science All Hands Meeting, September 2004, Nottingham, England.

Sinnott, R.O., Chadwick, D. W., Doherty, T., Martin D., Stell, A.J., Stewart, G., Su L., Watt J. (2008) *Advanced Security for Virtual Organizations: Exploring the Pros and Cons of Centralized vs Decentralized Security Models.* 8th IEEE International Symposium on Cluster Computing and the Grid (CCGrid 2008), May 2008, Lyon, France.

Snelling, D. and Kantarjiev, C. eds. (2007) *Technical Strategy for the Open Grid Forum 2007-2010* http://www.ogf.org/documents/GFD.113.pdf

Turner, A. (2008) *MoSeS-NCeSS e-Infrastructure use case,* http://www.geog.leeds.ac.uk/people/a.turner/organisations/OGC/GeoLinkingIE/MoSeS-NCeSS_e-Infrastructure_UseCase/