

Establishing a Persistent Interoperability Test-bed for European Geospatial Research

Gobe Hobona¹, Mike Jackson¹, Michael Gould², Chris Higgins³, Johannes Brauner⁴, Andreas Matheus⁵, Theodor Foerster⁶, Edward Nash⁷, Rob Lemmens⁶, Simon Abele⁸, Jerry Swan¹, Suchith Anand¹, Stefan Strobel⁵, Mohamed Bishr⁹, Peter Korduan⁷, Philip James⁸

¹Centre for Geospatial Science at the University of Nottingham, UK

²Universitat Jaume I, Spain

³Edina at the University of Edinburgh, UK

⁴Technische Universität Dresden, Germany

⁵Universität der Bundeswehr München, Germany

⁶International Institute for Geo-Information Science and Earth Observation, Netherlands

⁷Professorship for Geodesy and Geoinformatics, Universität Rostock, Germany

⁸School of Civil Engineering and Geosciences, Newcastle University, UK

⁹Universität Münster, Germany

INTRODUCTION

The development of standards for geospatial web services has been spearheaded by the Open Geospatial Consortium (OGC) - a group of over 370 private, public and academic organisations (OGC, 1999-2009). The OGC aims to facilitate interoperability between geospatial technologies through education, standards and other initiatives. The OGC Service Architecture, described in the international standard ISO 19119, offers an abstract specification for web services covering data dissemination, processing, portrayal, workflows and other areas. The development of specifications covering each of these categories of web services has led to a significant number of geospatial data and computational services available on the World Wide Web (the Web). A project¹ to establish a persistent geospatial interoperability test-bed (PTB) was commissioned in 2007 by the Association of Geographic Information Laboratories in Europe (AGILE), Commission 5 (Networks) of the European Spatial Data Research (EuroSDR) organisation and the OGC.

The PTB was established in recognition of the need to explore new and innovative approaches for facilitating interoperability between OGC web services. Further, the recent publishing of the INSPIRE directive acted as motivation for a European-oriented test-bed initiative (EC, 2007-2008). The PTB would stimulate collaboration in geospatial and SDI research. It would provide a mechanism for organisations to share data and services capability increasing the productivity of research by reducing duplication. It would also provide a benchmark for other consortia to compare their geospatial service-oriented architectures against. This paper introduces the PTB project and describes the use cases developed within Phase 1 of the project. A comparison to other test-bed initiatives is provided. Finally, a discussion of the limitations encountered during the development of Phase 1 is presented.

¹ <http://plone.itc.nl/gitest-bed>

SERVICE PLATFORM

A variety of specifications for geospatially-enabled web services have been published based on the ISO 19119, for example the Web Feature Service (WFS), Web Coverage Service (WCS), Web Processing Service (WPS) and Web Map Service (WMS). WFS offer an interface for disseminating vector data. WCS offer an interface for disseminating raster or coverage data. WPS offer an interface for publishing computational functions such as geometric algorithms or climate models. WMS offer an interface for publishing dynamically rendered maps or portrayals. Descriptions of resources offered by WFS, WCS and WPS can be retrieved by invoking the describeFeatureType, describeCoverage and describeProcess operations; the rest of this paper refers to them collectively as the describeResource operation. All OGC web services offer a getCapabilities operation that lists the resources offered by the services. By virtue of being based on specifications of the OGC, these services are collectively known as OGC web services. Table 1 presents a listing of mechanisms used by OGC web services for messaging, discovery and description.

	OGC
Messaging	OGC XML, ASCII Text (ArcGRID), GML, image binary (JPEG, TIFF etc)
Discovery	CSW, eBRIM
Description	describeResource , getCapabilities operations

Table 1: Listing of mechanisms for messaging, discovery and description in OGC web services.

IN CONTEXT

The OGC annually holds initiatives for testing new advances in open specifications. One such initiative is the series of OGC Web Services test-beds, named in sequence (i.e. OWS-1, OWS-2,..., OWS-6 and so on). Whereas early phases of the test-beds gave rise to specifications such as the WMS, recent phases have been instrumental in the development of geoprocessing and workflow specifications such as the WPS. The test-beds have also explored the integration of OGC web services with specifications from other standardisation bodies such as W3C and the Organization for the Advancement of Structured Information Standards (OASIS). Whereas the OWS series of test-beds is suitable for developing base interoperability specifications, there is still a need for a persistent Europe-oriented test-bed that can help with the implementation of the INSPIRE directive, GMES and other European initiatives. Further, the OWS series is primarily led by GIS vendors to test existing capability and explore near-term development options; therefore there is a need for a more research-oriented test-bed that can undertake more exploratory use cases and feedback findings into the longer-term OGC standardisation process.

As part of the establishment of a spatial data infrastructure (SDI) in Germany, a project was initiated called GDI NRW (Bernard, 2002). An SDI is a collection of technologies, people, policies and frameworks for facilitating the sharing of geospatial resources. The project was based on the state-wide sharing of geospatial data of various organisations through the use of OGC web services. The project involved the development of a test-bed for testing and developing interoperability mechanisms. Not only did the project allow participants to share data, but they were also able to share experiences and expertise. Such sharing of expertise and know-how is an important aspect of developing an SDI. Use cases within the project included scenarios in setup, publication, search,

ordering and use of OGC web services. In addition to well known services such as WFS and WMS, the test-bed also included a Web Pricing and Ordering Service (WPOS). The WPOS was one of the main contributions of the GDI NRW project. The project also highlighted the need for further research into how to integrate pricing and customer authentication into OGC web services through e-commerce techniques. The test-bed services were accessed through a forest management portal, developed as a prototype for the North Rhine-Westphalia Department for Forestry. The PTB is different to GDI NRW in that it includes more European countries and focuses on research collaboration. Further, the GDI NRW was developed before the publication of the INSPIRE directive. However, the GDI NRW is acknowledged as an important precursor to the PTB.

Another example of a geospatially-oriented test-bed is that of the Committee on Earth Observation Satellites (CEOS) - an international body that coordinates civil space-borne missions designed to observe and monitor the Earth (Di et al., 2003). In order to explore the possible application of grid computing in earth observation applications, CEOS established a test-bed that integrated grid computing and OGC web services. The test-bed included various use case demonstrators from NASA, NOAA, ESA and the USGS. The multi-organisation nature of the test-bed required a service architecture with a security layer to provide authentication and authorisation services. The OGC does not currently offer a security-oriented specification. Instead, the CEOS test-bed used the Grid Security Infrastructure (GSI) offered by the Globus Toolkit version 2.2. Using the Globus toolkit also provided a mechanism for resource sharing, one of the key benefits of grid technology. Geospatial capabilities were included through Grid-enabled WCS, CSW and WMS. As the CEOS missions offer data in HDF-EOS format, the WCS had the primary role of disseminating the data. For retrieval of rendered maps, the WMS was able to create maps from the HDF data. The CSW allowed clients to search and discover relevant data. The services have been packaged as the NASA Web GIS Software Suite (NWGISS). The integration of Grid computing and OGC web services has been explored in other studies such as the European Union-funded CYCLOPS² and SAW-GEO (Hobona et al., 2007).

The International Union of Geological Sciences (IUGS) developed an interoperability test-bed for exploring the exchange of data using the Geoscience Markup Language (GeoSciML) on the World Wide Web (Duffy et al., 2006). GeoSciML is an XML-based encoding for geoscientific data. It extends the Geography Markup Language (GML), a specification of the OGC. The project involved geological surveys from three continents, Australia, Europe and America. Initially the test-bed demonstrated dissemination of borehole data only, however the test-bed was later extended to support geological units, structures and other materials. The data was disseminated through web services compliant to OGC standards. The test-bed adopted three use cases; the first use case demonstrated the display, querying and retrieval of information in GeoSciML; the second use case demonstrated the query of multiple map features and download of a GeoSciML file; the third use case demonstrated the definition of various symbolisation schemes using GeoSciML attributes. The IUGS test-bed is an inter-continental initiative led by geological surveys and thus it is specifically for addressing needs of the global geoscience community. Consequently, it does not address Europe-specific issues of geospatial interoperability or provide a generic test-bed facility for collaborative GI research. However, it demonstrates the value of test-beds for advancing interoperability standards development within a specific geoscience community.

IMPLEMENTATION

At a series of meetings the PTB membership decided to develop four use cases for demonstrating geospatial interoperability, with the aim being to eventually expand the use cases to include more

² See <http://www.cyclops-project.eu/>

service providers. Consequently Phase 1 of the project involved pilot use cases on schematisation, secure portrayal, fog monitoring and semantic web services. In order to overcome funding constraints, the use cases reused elements of recent projects, consequently, a detailed description of the technical characteristics of the adopted services is provided in the referenced literature. The rest of this section describes each use case.




Figure 1: Authentication and Authorisation using GeoXACML, SAML and Shibboleth.

The first use case is based on secure portrayal of geospatial data. The use case was developed by the Universität der Bundeswehr München and EDINA, a national data centre based at the University of Edinburgh. It was implemented with the support of a consultancy called AM Consult. The basic idea was to develop some core PTB components using a hypothetical and generic cross border use case – the approach and the data could be reused in multiple scenarios. In this case, a scenario was explored where search and rescue teams from different countries have a need to portray maps with the same cartographic styles, e.g. in order to exaggerate particular features of interest, or simply to use familiar and consistent map portrayal. Secure use of individual countries cartographic representations was used as the vehicle for demonstrating access control to OGC web services - the project was resourced as part of the Joint Information System Committee (JISC) funded SEE-GEO project (Secure Access to Geospatial Web Services). The widely used Shibboleth system from the Internet2 consortium enabling browser based Single Sign On (SSO) across administrative boundaries was employed. Shibboleth is open source and based on open standards from OASIS, e.g. Security

Assertion Markup Language (SAML), XML encryption, XML signature. It is used for access management federations (primarily in the academic sector) in many countries including the UK, Germany, US, Finland, Australia. Basically, there are two main categories of components in Shibboleth access management federations; identity providers (IdPs) and service providers (SPs). In this case the protected services were OGC Web Map Services. A single IdP was set up by EDINA and two SPs in order to allow SSO to be demonstrated, one of the protected services was a WMS/WFS on top of pan European coverage of EuroGlobalMap – a 1:1 million data product from EuroGeographics. The authorisation mechanism used an implementation of the OGCs GeoXACML (Geospatial eXtensible Access Control Markup Language) specification. Figure 1 provides an overview of the architecture. The demonstrator allowed a user to authenticate using the conventional Shibboleth sequence of actions and then gain access to data suitably portrayed depending upon who they were, and what area they are requesting data for, e.g., Germany only, Netherlands only, cross border.

In summary, the key points to note are:

- The use case demonstrated an approach for protecting OGC Web Services without requiring changes to the OGC web service interface.
- The adopted approach requires use of a Shibboleth gatekeeper (or Policy Enforcement Point) around the OWS being protected.
- OWS clients have to be modified to work with Shibboleth.
- The use case opens up the possibility of a PTB Access Management Federation.
- The pan-European data services may be reusable with appropriate licensing constraints.

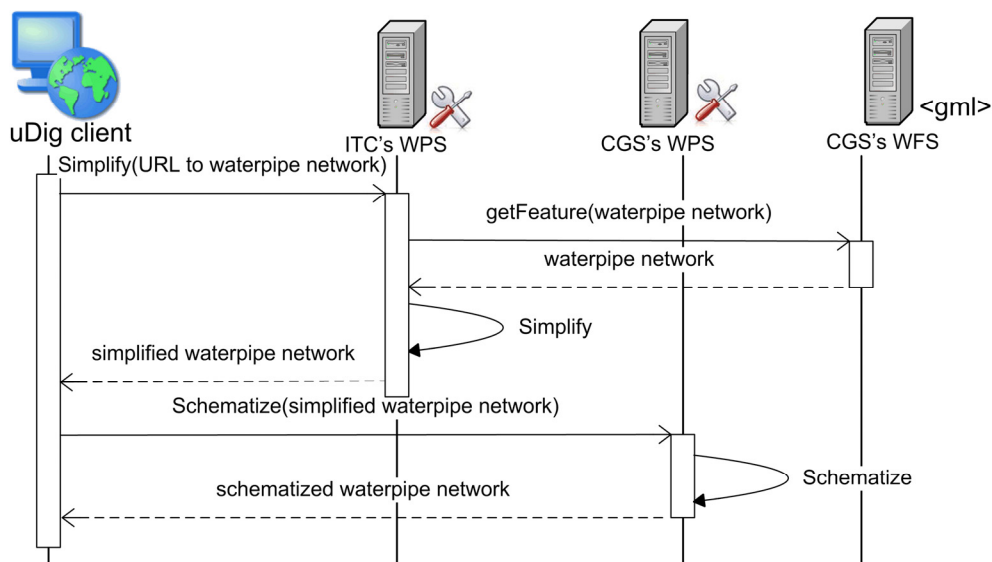


Figure 2: Schematisation Use Case Sequence Diagram.

The second use case is based on schematisation, a method of modifying geometries to make a network map easier to read (the classic example of which is the map of the London underground). The Schematisation use case was developed by researchers from the Centre for Geospatial Science (CGS) at the University of Nottingham and the International Institute for Geo-Information Science and Earth Observation (ITC) in the Netherlands (Swan et al., 2008). The use case involved WFS and WPS instances hosted in Nottingham (United Kingdom) and Enschede (Netherlands). The use case is based on a water utility scenario triggered by a faulty pipe in the utility's network. A maintenance engineer is dispatched to attend to the faulty pipe thereby requiring both a topographic map and a schematised map showing network topology. The engineer identifies relevant resources through a Catalogue Service for the Web (CSW). Through a desktop GIS such as uDig³, the engineer activates a workflow that retrieves network geometry data from the WFS instance in Nottingham. The network geometry is then sent to the WPS instance in Enschede which simplifies the geometries to reduce the number of vertices in the network without significantly deforming the geometries. The simplified geometries are then forwarded to the WPS instance in Nottingham that applies a schematisation algorithm on the geometries.

The software packages applied for this use case are available under Open Source license (geoserver as WFS implementation & 52°North WPS as WPS implementation). The water pipes data are encoded as GML. Moreover, the geospatial data exchanged between the different WPS instances are not directly included in the messages, but included as URL references. Using references to geospatial data reduces the communication overhead and enables possible server-side caching of already received data. The communication between the different services to realize the Schematization use case is depicted in Figure 2. The processing of the 200 edges of water pipe data took approximately two minutes. This included the time taken for message transfer between the services and the time for geometrically processing the data. The relatively long response time of the service chain shows, that performance is still an unsolved issue in this use case and will be tackled in the future. However, the use case shows that in a fairly short period of time a distributed process can be setup by different service authors with complementary functionality based on the WPS interface specification.

³ <http://udig.refractions.net>




Figure 3: Component Communication in Fog Information Use Case.

The third use case is based on a web-based information service and is developed by Newcastle University and Universität Rostock. The use case is centred on the use of a light-weight Client for OGC Web Services (COWS, described in Nash et al., 2008) for accessing a variety of data sources and publishing alerts to interested parties. The COWS aim is to use Asynchronous Javascript And XML (AJAX) techniques to produce a flexible and generic browser-based client for all OGC services. Such a client should play a central role in the test-bed by providing a tool for ad-hoc inspection and chaining of contributed services. The scenario presents a transport engineer tasked with creating a map of fog hotspots based on data retrieved from a sensor network. Through a COWS client the transport engineer obtains a density coverage based on observations from fog sensors. The density coverage is created through a WPS process that interpolates between observations retrieved from SOS. Fog hotspots are then identified through raster subtraction of the density coverage from a threshold coverage. The threshold coverage is retrieved from a WCS. The transport engineer is then able to digitise around the fog hotspots using a transactional WFS, thereby enabling motorists to view the fog hotspots via a WMS using a mobile client. An illustration of communication between components in the use case is presented in Figure 3.

From this use-case, the WPS services (interpolation and map algebra) are pre-existing as they have been developed within other projects, although only conform to the draft 0.4.0 WPS standard and not version 1.0.0. These services provide fundamental geoprocessing functions which it is expected will be widely used in other use-cases in future. Similarly, SOS services with weather data are also available, providing sample sensor data which may be relevant in other use-cases. The direct implementation of the COWS client has stalled due to funding constraints and personnel changes, however much of the functionality which was initially envisaged is currently in the process of being implemented within existing AJAX frameworks (e.g. OpenLayers).




Figure 4: Overview of Semantic Discovery Environment (Source: Andrei et al. 2008).

The fourth use case applies semantics and ontologies to OGC web services, using SWING technology (Andrei et al., 2008). The use case is themed on a geologist requesting a production-consumption map of aggregates for a construction project. Datasets included in the map include quarry locations, production at quarries, administrative boundaries and population (census) data. All the datasets are provided by WFS. The publisher enters a Knowledge Base (KB) -a collection of concepts and assertions binding the concepts. The publisher registers the service as being applicable in a production-consumption map scenario. A geologist accesses the KB to search for an aggregates consumption map. A map created with services annotated with concepts from aggregates ontology is retrieved. The geologist is then able to edit the map by setting the presentation styles of the aggregates, administrative boundary and population density datasets. An illustration of the use case is presented in Figure 4.

DISCUSSION

As the number of services involved in the PTB increases, the need for access control is expected to increase. Security provision has not been adequately addressed by OGC web services; however the work undertaken in the first use case offers the possibility of future projects including Web Service Security (WS-Security) within geospatial service frameworks. With the institutional membership of the PTB, it is evident that a federated authentication model is required, in contrast to having a single authentication authority. The participation of 25 institutions means that managing authentication for the complete project could require additional human and financial resources that are currently not available. Therefore, a federated approach, possibly adopting the Shibboleth-based approach developed by EDINA, could provide a solution for the PTB and similar collaborative test-beds.

Regarding authorisation, it can be expected that multiple roles will be required to address copyright constraints that vary with each country. For example, upon signing an Ordnance Survey agreement, UK-based researchers could be allowed to access data services that hold Ordnance Survey data. However, to preserve the community of the PTB, some resources will need to be made universally accessible.

A key concern for the PTB will be to demonstrate its utility through the addition of more services than those implemented in the pilot use cases. As additional services are included in service chains, the complexity of the processes will require orchestration through workflow enactors. This may result in the PTB including a repository of workflow templates. Repositories of workflow templates are already gaining popularity in the bioinformatics community through the MyExperiment.org project. A critical mass will likely be reached when the PTB is of sufficient maturity to allow creation and storage of geospatial workflows, thereby requiring research into the capacity of CSW and geospatial metadata to support discovery of geospatial workflows. It is envisioned that until a mechanism is provided within the PTB for workflow enactment, users will invoke services through standalone applications or web-based tools as demonstrated by the second and third use cases.

A major challenge for future development of the PTB will be to build on the enthusiasm of those 25 organisations that responded to the original call for participation in the Summer of 2007. Of particular note is that several of those organisations volunteered OGC compliant catalogue services adhering to various profiles. This reflects the important role that resource discovery has within Spatial Data Infrastructures and the need for extensive related research and development. The second and fourth use cases both adopt catalogue services within their workflows. The central role of catalogue services will become more apparent as the number and variety of resources need registering as participating in the PTB increases. It may be desirable (and practical) to adopt a federated approach to catalogues from the outset. A related issue of equal importance will be the establishment of appropriate governance mechanisms to maintain coherence amongst these diverse resources.

A key result of this project has been the development of a template for a persistent test-bed that could be used to explore interoperability between components of a regional SDI. This has led to the involvement of the EU FP7 GIGAS forum and EU eContentPlus ESDIN project in the coordination of the PTB. GIGAS aims to facilitate interoperability between GEOSS, INSPIRE and GMES initiatives. ESDIN aims to facilitate the development of a European SDI through improving the ability of SDIs of member states to share data and services. The multi-national nature of the institutions participating in the PTB and the fact that the contributed services are hosted in different countries should provide valuable feedback to both the GIGAS⁴ and ESDIN⁵ projects.

CONCLUSIONS AND FUTURE WORK

This paper has presented a strategic project to develop a persistent interoperability test-bed for geospatial research in Europe. The project, though initially unfunded, has been supported through the contribution of web services and technologies by academic geospatial researchers from across Europe. Implementation work in the project involved adapting research products to specified use cases, thereby overcoming funding limitations. We conclude that an access control model that supports federated authentication may be required to support the multitude of institutions that have shown an interest in the PTB. Further, we conclude that a repository or registry of workflows will be required as more services are added to the use cases. It is envisioned that future work on the project

⁴ <http://www.thegigasforum.eu>

⁵ <http://www.esdin.eu>

will provide closer integration with the development of European initiatives such as INSPIRE and GMES.

ACKNOWLEDGEMENTS

The authors acknowledge the support of AGILE, EuroSDR, OGC, the GIGAS forum and ESDIN.

BIBLIOGRAPHY

- ANDREI, M., BERRE, A., COSTA, L., DUCHESNE, P., FITZNER, D., GRGAR, M., HOFFMANN, J., KLIEN, E., LANGLOIS, J., LIMYR, A., MAUE, P., SCHADE, S., STEINMETZ, N., TERTRE, F., VASILIU, L., ZAHARIA, R. & ZASTAVNI, N. (2008) SWING: An Integrated Environment for Geospatial Semantic Web Services *Lecture Notes In Computer Science*, 5021, 767-771.
- BERNARD, L. (2002) Experiences from an implementation Test-bed to set up a national SDI. IN RUIZ, M., GOULD, M. & RAMON, J. (Eds.) *5th AGILE Conference on Geographic Information Science*. Palma de Mallorca, Spain.
- DI, L., CHEN, A., YANG, W. & ZHAO, P. (2003) The Integration of Grid Technology with OGC Web Services (OWS) in NWGIS for NASA EOS Data. *Global Grid Forum(GGF) 8 & High Performance Distributed Computing(HPDC) 12*. Seattle, Washington.
- DUFFY, T. R., BOISVERT, E., COX, S., JOHNSON, B. R., RAYMOND, O., RICHARD, S. M., ROBIDA, F., SERRANO, J. J., SIMONS, B. & STOLEN, L. K. (2006) The IUGS-CGI international geoscience information interoperability test-bed. *XIth International Congress of the International Association for Mathematical Geology*. Belgium.
- EC (2007-2008) INSPIRE Directive and Data Models. European Commission Joint Research Centre (JRC).
- HOBONA, G., FAIRBAIRN, D. & JAMES, P. (2007) Workflow Enactment of Grid-Enabled Geospatial Web Services. *Proceedings of the 6th UK e-Science All Hands Meeting*. Nottingham.
- LEMMENS, R., WYTZISK, A., BY, R. D., GRANELL, C., GOULD, M. & VAN OOSTEROM, P. (2006) Integrating Semantic and Syntactic Descriptions to Chain Geographic Services. *IEEE Internet Computing*, 10, 42-52.
- NASH, E., KORDUAN, P., ABELE, S. & HOBONA, G. (2008) Design Requirements for an AJAX and Web-Service Based Generic Internet GIS Client. *11th AGILE Conference on Geographic Information Science*. Girona, Spain.
- OGC (1999-2009) Open Geospatial Consortium specifications.
- SWAN, J., FOERSTER, T., LEMMENS, R., HOBONA, G., ANAND, S. & JACKSON, M. (2008) Discovery and Invocation Of Schematization Services: A Use Case For OGC-EuroSDR-AGILE Persistent Test Bed For Europe. *GEOconnexion International magazine*, 7, 24-27.