

Grid and OGC Collision

This document describes a proposed JISC call for work linking together web service activities in the geospatial community and the UK national e-science community.

Motivation and Background

1. The Open Geospatial Consortium, Inc. (OGC) is a non-profit international voluntary consensus standards organization that is leading the development of standards for geospatial and location based services. OGC activities are especially relevant to JISC in that geospatial services are inherent in both national infrastructure participation and research services that JISC member bodies need to provide.
2. OGC consensus leads to OGC specifications that are divided into Abstract Specifications and Implementation Specifications.
3. The Abstract Specifications provide the conceptual foundation for most OGC specification development activities, and are being harmonised with emerging international geographical standards being generated by the International Organisation for Standardisation (ISO). Open interfaces and protocols are built and referenced against the Abstract Specification, thus enabling interoperability between different brands and different kinds of spatial processing systems. The Abstract Specification provides a reference model for the development of Implementation Specifications which target a technical audience and detail the structure of the interfaces between the distributed software components.
4. The Grid as defined by Foster (<http://www.gridtoday.com/02/0722/100136.html>) has three main characteristics:
 - standard open protocols,
 - no centralised control, and
 - non-trivial quality of service.The lack of central control requires secure access to grid resources, typically implemented through public key infrastructures.
5. Existing and planned OGC web services are open protocols designed for the situation where there is no centralised control, and are intended to provide non-trivial quality of service.
6. The OGC has been developing Web Service interfaces over a number of years, albeit based on RESTful principles (before the term was invented). One of the major identifying characteristics of OGC web services is that will provide identifiers to datasets and services identified using the citation element of standardised metadata constructs (ISO19115).
7. Currently there are a number of published standard web service interfaces, and a number of activities that are in various states of maturity. The flagship OGC standard web services include the
 - Web Coverage Service, which returns an array of numbers from a specific product covering a specific geospatial domain
 - Web Map Service, which returns an image map of a specific product on a specific domain.

- Web Feature Service, which returns a specific geolocated feature of geographic interest. (Features should themselves be described and registered using appropriate services).
8. In general OGC web services are based on existing International Standards Organisation standards for geospatial information, which include methodologies which span the gamut from identifying the problem, via UML conceptual modelling and standard techniques for XML encoding, to producing both open reference implementations and closed commercial implementations of the standards.
 9. In recent years, the OGC has commissioned work on producing SOAP-based implementations of the existing OGC web services as well as further developing the suite to include processing services. Currently the OGC is refining best practice in geospatial service architectures for
 - Publish, find and bind
 - Push and pull publishing (harvesting)
 - Service description
 - Service taxonomy
 - Distributed searching
 - Semantic data integration
 10. In practice then, the geospatial community may regard the development of “grid” principles as essentially the re-development and wider deployment of concepts already prevalent in their community. However, there are significant areas where the grid community is more advanced, and using grid concepts could significantly improve the capability of the geospatial community to exploit existing datasets, and hence improve research and educational productivity for the JISC community.

Security

11. With the exception of a recent discussion document (Matheus, 2005¹) and the newly formed Digital Rights Working Group (GeoDRM) the OGC community has done little work on secure access to services, relying on guaranteeing secure access by controlling the access layer. However, both Lawrence *et al* (2004²) and Matheus have argued that methodologies based on controlling access to individual data elements behind secure services are necessary for scalable and interoperable services across differing management domains.
12. By contrast, the Grid community has invested significant resources into both authentication (based on proxy certificates and PKI) and authorisation (in a number of projects, including, but not limited to PERMIS³, AKENTI⁴ and the NERC DataGrid as described in Lawrence *op cit*). While there are two possible core web service architectures for grids likely to be prevalent in the near future (WSI and WSRF, see Fox 2005⁵), both share usage of WS-Security for message level security. Grid authorisation protocols based on standard SAML and XACML would also be the preferred technologies for OGC deployment.

¹ GeoXACML, a spatial extension to XACML. OGC discussion paper 05-036, https://portal.opengeospatial.org/files/index.php?artifact_id=10471

² Lawrence, B.N. *et al.*, The NERC DataGrid: Googling Secure Data, <http://www.allhands.org.uk/2004/proceedings/papers/76.pdf>

³ Privilege and Role Management Infrastructure Standards Validation, see <http://sec.cs.kent.ac.uk/openPermis/>

⁴ <http://dsd.lbl.gov/Akenti/>

⁵ Geoffrey Fox: Possible Architectural Principles for OGSA-UK and other Grids, http://www.nesc.ac.uk/events/townmeeting0105/talks/ogsaukjan05_Geoffrey_Fox.ppt

Architecture and Workflow

13. OGC web services can, or are intended to, participate in geoprocessing chains including, for example, the Web Processing Service (WPS), or demonstrated by applications which exploit multiple integrated products, for example cascaded maps from various web map servers. The Web Processing Service concept includes support of
 - a) integrated products (as for example layered maps from OGC web map servers),
 - b) major data requests including notification, state maintenance, and marshalling of resources
 - c) multiple providers.
14. ISO19119 is an abstract description of how geographic information services can be described and categorised within a service taxonomy, and tested as conformant to the categories. ISO19119 service descriptions are intended to be implemented in the OGC catalogue services as service registries. Amongst other things, ISO19119 identifies the importance of service chaining and describes workflow services including “chain definition service” and “workflow enactment service”
15. The UK National e-science programme has put significant work into developing and implementing workflow technologies (see for example the list of projects discussed at the 2003 meeting on work flow services (<http://www.nesc.ac.uk/esi/events/303/>)).
16. In general existing grid workflow technologies provide the capabilities of “chain definition service” and “workflow enactment service” but would need to be modified or enhanced to support the use of OGC web services (including the catalogue services which provide registries). Note that Fox, *op cit*, has suggested that architectures should try and separate application services from core services, but that OGC services will be difficult in that they involve system services (i.e. the catalogue and processing services).
17. OGC services are intended to provide access to datasets and services described by OGC catalogue services and ISO19115 citations. At the basic level, it can be argued that this paradigm might not best be supported by web services that only provide access to datasets mediated by specific hosting services. Grid services which exploit WS-Addressing or similar technologies might provide access to data entities independently allowing the construction of more complex workflow.

Prior Activities

18. An ad hoc GridGIS working group was established followed a meeting held under the auspices of the NERC e-science programme (<http://www.niees.ac.uk/events/GridGIS/>)
19. JISC funded an interoperability project (<http://edina.ac.uk/projects/interoperability>) which recommended that

“a specific working group (perhaps the GridGIS Working Group) be tasked with investigating web services security issues and that workshops (incorporating the ongoing JISC funded work on Shibboleth) be convened in order to establish best practice and to recommend a road map for addressing security and geoDRM aspects of geospatial interoperability within the UK academic community.”

20. JISC JCSR will be (if JCSR supports it) funding a demonstration pilot to investigate exploiting existing grid technologies (explicitly, the National Grid Service and the GSI security infrastructure) in supporting the work of EDINA and MIMAS. This call will differ significantly in that it concentrates on next generation grid services, may or may not involve existing JISC services, and on explicitly engaging with the OGC on taking up grid tooling, to be contrasted with explicitly tasking existing JISC data services to exploit existing infrastructure (as planned in the other activity).
21. The OGC has an Earth Observation Working Group which has focused recently on Grid topics, and in partnership with the Committee for Earth Observation Satellites (CEOS) is sponsoring test bed activities, see <http://wgiss.ceos.org/meetings/wgiss20/Tech & Servs/GRID/>
22. The OGC Web Service-3 Initiative (<http://www.opengeospatial.org/initiatives/?iid=162>) has recently supported work in a variety of areas of which the two most relevant to this JISC call are the common architecture and geographical digital rights management. These two areas include recognition that the basic common OGC web service architecture may need enhancement/modification to support more complex interactions.

Statement of Work

JISC will receive calls in three areas:

- I. Building prototypes of OGC services that exploit WS-Security and other appropriate grid based security mechanisms. An important requirement of successful prototypes will be the ability for servers to provide secure access to resources at the dataset level, not just via secure wrapping of the service delivery. It is expected that a successful response to this part of the call would include work to produce:
 - a. A secure prototype of one or more OGC web services (with a demonstration client), and if necessary/appropriate:
 - b. Proposals to OGC on how to modify the OGC web service standards to support such secure access (such proposals will have to use OGC compliant architectural terminology).(Maximum funding: £100,000. Maximum duration: two years).
 - II. Developing and extending workflow engines to support and orchestrate OGC services (including prototypical examples of WPS wrappers to other grid services) described using ISO19119 and understanding the relationship between ISO19119 and other grid service registries (to possibly include engineering ISO19119 compliant versions of existing grid service registries).
(Maximum funding £100,000. Maximum duration: two years).
 - III. Attendance at and contributions to OGC technical meetings and reporting on developments to the e-science community via written reports and presentations.
(Maximum funding: £10,000 per annum for two years).
23. Proposals will be expected to clearly identify
- o Which one (or ones) of these areas for which support is sought.
 - o How they will contribute grid technologies to solving the problems involved in interoperable geospatial data access.

- Why they should be funded by the JISC as opposed to be directly involved in the OGC activities underway.
- Although recognising that there are differences between JISC and OGC, proposals should also demonstrate how far they fit into:
 - i. The overall service orientated JISC architecture, and
 - ii. The international standards based OGC architectural philosophy.

24. It is expected that, where possible, proposed activities will exploit interactions with existing OGC activities (e.g. the OGC interoperability programme). Note however, that JISC will not fund OGC membership to expedite such interactions, which we assume is funded via other mechanisms.