

A03.1

Web Services: the key at last to “joined-up geography”?

Ed Parsons, Chief Technology Officer, Ordnance Survey

Abstract

The IT industry is currently calling Web Services the “next big thing” and it is clearly the recipient of some considerable hype at this point in time, nevertheless the concept of intelligent system to system communication may solve many of the problems of joining up disparate geospatial data sets and services.

Web services are the building blocks of the future distributed computing platform on the Internet. Open standards here are key, with a focus on communication and collaboration among people and applications rather than strict descriptions of particular geographical data models. This is possible through the use of XML schemas such as GML as the mechanism for application integration.

This paper will investigate how through the use of Web Services and the technologies of SOAP and UDDI organizations such as the National Mapping Agencies, Government departments, Software Vendors and System integrators can interface their data and services to provide seamless applications which leverage their combined offerings transparently to the end user.

For example a future web services user may login to a building society branded website which has an online housing buying assistant application. The application itself a web service could communicate with other web services across the internet to identify if a potential property, was a risk of flooding, had a registered title, could be connected to cable TV etc. The user might also register to have mail redirected, utilities and services connected and the electoral role updated. Each of the organization systems involved would remain dedicated to that organisations operation but would present a web services compliant interface to the others to allow meaningful collaboration.

Such integration and collaboration has often been promised in the past, but now finally we have a technology which can begin to deliver joined up geography.

The problems in joining geographies

The power of geography as an integrating force is well recognised allowing disparate sets of data to be brought together based upon a common location so that the underlying relationship between for example poor health and industrial pollution can be identified. This can only take place if both data sets have been collected on a common framework or geography which allows their location to be identified, the use of coordinates such as the National Grid reference or Latitude and longitude is the most common direct method. In most cases however the location must be identified using another geography from which we can indirectly imply a location, Postal addresses, UPRN, TOID, census output areas, and parliamentary constituencies are all examples of these alternative geographies.

The geocoding process of obtaining a discrete location from these many geographies is developing into quite a cottage industry with much effort been duplicated across the GI industry in the UK to bring these datasets together. The conflagration process often involves both an understanding of the underlying data structures and formats but also the application of geographical information science in combining datasets with different temporal and spatial resolutions.

The largest problem in joining geographies is however organisational, identifying who has information, how the information was collected and how it is structured (metadata) and actually getting access to the data itself on CD, Tape paper or even online !

To use an analogy to describe access of geographical information in the UK today, we are currently experiencing the early days of the internet when to read an article or document on a server on the internet, a user would need to know the IP address of the actual server, the access protocol to download the document (who remembers gopher?) a binary encoding/decoding application, the file format of the document and the appropriate file viewer. Today of course we just point our browser at google.com, enter some relevant text as a search string and view the pdf document! We are still some way of bringing that level of integration to the proprietary and specialised realm of GI data but the technology of web services appears to hold the key to this level of usage.

Web services defined

So far the internet and the World Wide Web developed on top of it, has been about the dissemination of digital information from a server computer somewhere to a user using a client computer somewhere else. The information is transported and presented in such a way that a human is really required to interpret the information. The technology behind Web services removes the need for human interaction allowing different software applications to communicate and share information in a well defined and structured way allowing processing to take place as a “back office” activity. For this reason despite all the current hype surrounding web services many users of the internet or IT systems in general may not notice the revolution taking place behind the scenes of their familiar application.

A more technical definition of a web service is a “URL-addressable software resource that performs functions and provides answers” (Seybold, 2002). The Web Service is an encapsulation of perhaps existing software functionality in a common form that allows the services it performs to be visible and accessible to other software applications. A single web services based application can request services from other Web Services, and can expect to receive the results or responses from those requests in an expected form.

The great advantage over existing integration technologies is that Web Services are designed to interoperate in a loosely-coupled manner; they can request a particular type of services across the internet and wait for responses. A Web Service can be discovered and used by other Web Services, applications, clients, or agents. Web Services may be combined or chained to create new services. And they may be recombined, swapped or substituted, or replaced at runtime.

To illustrate this, an organisation might require as part of their corporate portal a location map, which could be delivered as a web services from a map publisher which in turn obtains the raw spatial data from a street centreline database developer which in turn receives data from a web service publishing data from the national mapping agency. Any one of these providers could be changed to another in real time without any application development or change in functionality. This change might occur because the data feed might be disrupted, or a “better” supplier might be identified as the preferred choice.

The mechanics of GI web services

A web service is “published” to the network inside or outside the corporate firewall by providing a document describing its functionality and mode of operation. This document is created using the core building block of web services XML the eXtensible Markup Language, an “information rich” description language. For this particular purpose a particular type or schema of XML is used to create the document. The web service is described using the Web Services Definition Language (WSDL), of particular note here is that the WSDL document is machine readable so that the capabilities of the service can be explored by another web service. In the context of GI web services a WSDL document might describe a service which when passed a national grid reference responds with the relevant LandRanger map and its details. The WSDL describes how the service is evoked and what is returned.

To make the rest of the network aware of this web service the WSDL is published to a public clearinghouse or yellow pages site known as a Universal Discovery, Description, and Integration or UDDI registry. A developer of web services or a web service itself can therefore either couple with another known web service or may use UDDI to discover equivalent services.

The communication between web services is carried out by passing XML messages wrapped in an interoperable container to allow the messages to cross different networks, use different application architectures and systems. The wrapping mechanism is known as SOAP the Simple Object Access Protocol and is best thought of as a standard DL sized envelope which keeps the contents of the message secure and also guarantees the message will be deliverable through a standard letterbox.

Potential GI web Services

So how might the concept of Web Services and the technologies of WSDL, UDDI and SOAP combine to deliver benefit to the GI industry? From the perspective of a National Mapping agency, the Ordnance Survey can see a number of valuable Web Services which could be developed internally and in conjunction with our partners which would aid in the uptake of Geographic Information largely by removing much of today's complexity.

A Mapping web service

The Ordnance Survey like many organisations is making increasing use internally and externally of web based mapping applications which generate on demand maps based on our existing cartographic products. These applications could be encapsulated as web services so that a simple SOAP message could evoke them and generate a map. The building blocks for such applications are already mature and much of the work of the Open GIS consortium has been directed at developing the protocols for such web mapping services (OpenGIS, 1999). As with all web services the sophistication of simple web service based application may be extended by cascading or chaining services together so that a number of web map services may combine their maps and then be published through a transaction based commerce server (Wagner, 2002)

A "Change Only" update service

One of the most powerful aspects of the new MasterMap Topo data product is the ability to supply only updated topographic features online rather than complete resupply on update. However with this advance it must be recognised come higher data management requirements for the user of change only update data.

A series of web services could be developed which would allow a user organisation to subscribe to an update service which would message the changed features to a receiving service which could then automatically insert the changes into the users own local repository. From the perspective of the local system administration this could offer up to date consistent data with very little management overhead.

This same web service application could also be repurposed to serve the Ordnance Survey's value added partner community allowing them to keep their own local databases synchronised with the National Topographic Database.

A Web Feature service

Building on from the previous two services, a more sophisticated web service could deliver raw Geospatial feature data to end users or partners for local rendering into maps or processing. With the appropriate support infrastructure a web feature service could deliver spatial data directly into a user's desktop GIS application removing the requirement to manage any local repository of data. Such a feature service could also be chained through a Geospatial analysis web service which could allow analysis of the data to be carried out on the data as it passes through, for example calculating feature centre points or overlaying features from different geographies.

The chaining of data and Geospatial applications could be vital in the delivery of successful Location Based Services where the end users terminal has limited processing and display capabilities. The "Find the nearest Indian restaurant" example could be delivered by chaining web services publishing Points of Interest, Integrated Transport Networks, providing Routing, Map rendering, and text to voice conversion. In the competitive mobile communications marketplace it might be that the operator with the best collection of branded web service providers becomes the dominant one.

Miscellaneous services

Because the GI web services in most cases will have no direct user interface and will become embedded deeply within other systems, many of the web services developed to support the GI industry will be invisible to the end user. These small services however could be some of the most powerful; a gazetteer or place

name service could be developed along with tools for identifying locations based upon postcode or census geographies. An important service which could be developed could provide lookups between the property reference number of the National Land and Property Gazetteer and the OS MasterMap TOID.

Some issues remain

As with an emerging technology there are risks in embracing web services at this point. Although a new technology with immature tools the development environment for the creation of web services is relatively straightforward. There are competing technology platforms in the form of .NET and J2EE from Microsoft and Sun respectively, but there is complete interoperability between them.

By far the biggest challenges are cultural rather than technological. From the perspective of a system developer the move to a web services based architecture is a radical change from the closely coupled largely proprietary world of GIS. A vendor will no longer be able to “own” the customer by providing a closed system, in fact with web services the system may only come together at run time for a single instance, and never be formed of those exact components again.

From the perspective of a data supplier greater interoperability means more focus on issues of accuracy, completeness currency etc which in a less interoperable world were hidden.

Perhaps the biggest area of concern is the immaturity of business models on the internet in general but more specifically the rather more complex (from a business workflow perspective) world of web services. I believe there is still considerable work required to develop a viable business model robust enough to cope with a service delivered onto a mobile user’s terminal for the cost of a few pence which is the product of a complex chain of web services involving multiple data and application service providers.

Conclusion

The GI industry and to a lesser extent the internet are technologies which are not fulfilling their obvious potential. The potential for change that web services bring to increase the actual usefulness of the internet is massive. Web sites will become truly interactive, able to respond to the user’s needs, and the environment using application logic. These same advances offer the potential to increase the use of Geographic information by removing some of the complexity and reducing the high cost of entry. Users of GI data be they professional planners or special interest activists will be able to gain access to the data they need and the tools to use it when they need and only when they need it.

Joined up geography will become less of a issues in the web services landscape as all data will appear to be joined up, as behind the scenes web services work 24/7 to manage and manipulate the data in a user rather than system centric view of information.

The Ordnance Survey is continuing to investigate the potential benefits of web services, to ourselves, our partners and our customers. Please feel free to contact us with your needs/issues as no organisation can hope to develop web services in a vacuum.

Seybold, P.A. 2002 “Web Services Guide for Customer-Centric Executives”, Patricia Seybold Group, Inc. Boston

OpenGIS Consortium 1999 “Open GIS Consortium Web Mapping Testbed”, www.opengis.org/archive/wmt

Wagner, Roland M. 2002 “GIS Meets E-Business: Web Pricing and Ordering Service (WPOS)”, Proceedings on the 25th GITA Annual Conference & Exhibition