Open Geospatial Consortium

Date: 2011-04-07

Reference number of this OpenGIS© Project Document: **OGC 11-036**

Category: OGC© White Paper

Editor: Lance McKee, Carl Reed, Steven Ramage

OGC Standards and Cloud Computing

Warning

This is a draft OGC White Paper. The document is available for comment.

Document type: OGC® White Paper
Document subtype: N.A.
Document stage: Draft
Document language: English

Participating in the OGC (see http://www.opengeospatial.org/contact) provides unique opportunities for organizations to track and influence geospatial technology and gain early understanding of related business issues. There is no better way to become familiar with the OGC Web Services standards that enable value to be derived from geospatial cloud computing solutions.

FORWARD

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. The Open Geospatial Consortium Inc. shall not be held responsible for identifying any or all such patent rights.

Recipients of this document are requested to submit, with their comments, notification of any relevant patent claims or other intellectual property rights of which they may be aware that might be infringed by any implementation of the standard set forth in this document, and to provide supporting documentation.

Acknowledgements

The OGC would like to thank the following individuals for their comments and suggestions regarding the content of this white paper.

Abstract

This OGC White Paper discusses cloud computing from the perspective of OGC's geospatial standards development activities and standards baseline. The paper begins with a discussion of what the cloud and cloud computing are. Unfortunately, there is still considerable misunderstanding in the geospatial technology community regarding cloud computing. The paper then discusses how standards figure into the options, benefits and risks of cloud computing for users and providers of geospatial data and software. This perspective is important not only for those immersed in geospatial technology, but also for cloud service providers, customers and technology partners who may be unfamiliar with the basic issues surrounding geospatial technology. This white paper does not discuss vendor specific cloud computing platforms.

What is the Cloud and Cloud Computing?

The following is a discussion the cloud and cloud computing. These definitions are important as there is considerable misunderstanding in the geospatial community as to what is cloud computing. In an October 2010 survey (unscientific) Directions magazine asked users about their understanding of cloud computing. About 25% of the respondents had ever heard of cloud computing and another 30% were not sure what cloud computing is. Even more interesting is that among those that believe they know what cloud computing is, there is still considerable misunderstanding. For example, one author suggested that a single quad-core server running GIS software and connected to many clients via CITRIX was a scalable cloud computing platform. This is not correct and again suggests that there is considerable misunderstanding in the geospatial community about the cloud and cloud computing.

The cloud concept is not new. For decades, presentations have used a "cloud" icon to depict the internet and/or the web. To users, the internet cloud is an amorphous infrastructure of computers, networks, and software. Most users have no idea how the internet or the web works. Beyond hardware and networks, standards enabled and continue to enable the evolution of the internet infrastructure and the world wide web. Without standards, there would not be the internet or the web. The very first standard developed to enable the web was TCP/IP. In the same context, the value of standards in the growth and evolution of the cloud and cloud computing is an important concept to keep in mind when geospatial cloud computing is discussed.

Another key aspect the cloud is client-server. Client-server applications utilize a distributed architecture, such as the internet, in which there are providers of a resource or service, called servers, and service requesters, called clients. Typically clients and servers communicate over a network on separate hardware, but both client and server may reside in the same system. A server machine is a host that is running one or more server programs. A client requests a server's content or service function. Clients therefore initiate communication sessions with servers which await incoming requests. This is the definition of client-server as used by the OGC community.

As with the "cloud", client-server applications have been used since the 1960's. For example, beginning in 1965 Dartmouth College had a GE timeshare system with remote terminal access via phone lines (300 baud!). However, not until the mid 1990's and the first web browsers did client-server evolve into the normal mode of operations for the vast majority of users, especially for geospatial content. With the advent of the web, clients such as Mosaic or Netscape offered users access to applications on servers. No user really knew or cared where the servers were located. They were accessible. As long as users had internet access, they had access to these applications. Again, standards such as HTML, HTTP, and URIs allowed this major evolutionary leap in the development of the cloud to happen.

Wikipedia provides a definition of cloud computing as, "Cloud computing is Web-based processing, whereby shared resources, software, and information are provided to computers and other devices (such as smartphones) on demand over the internet". By this definition, we have been doing cloud computing since the mid-1990's!

So what is cloud computing in 2010? The National Institute of Standards and Technology (NIST) was tasked with developing a common definition for cloud computing[1]

From that document:

Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model promotes availability and is composed of five essential ***characteristics***.

A synopsis of these five essential characteristics are:

On-demand self-service. A consumer can unilaterally provision computing capabilities without requiring human interaction with each service's provider.
Broad network access. Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, laptops, and PDAs).
Resource pooling. Computing resources are pooled to serve multiple consumers with resources dynamically assigned and reassigned according to consumer demand. The customer generally has no control or knowledge over the exact location of the provided resources.
Rapid elasticity. Capabilities can be rapidly and elastically provisioned.
Measured Service. Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service.

The reader is encouraged to read the entire NIST definition for cloud computing. A key aspect of cloud computing are pooling and elasticity. These two characteristics are especially relevant to geospatial cloud computing.

The Cloud , value and economies of scale – The cloud today

As has been discussed, cloud computing in 2010 is the result of a natural evolution and integration of the internet, the web, web services, and the very old client server model. At the end of the day, the cloud and cloud computing are the result of economies of scale

[1] http://csrc.nist.gov/groups/SNS/cloud-computing/cloud-def-v15.doc

applied to the provisioning of computing resources in a world of broadband Internet, "virtualization", open Internet and Web standards, and demand for an ever increasing number of on-line web services. To a considerable extent, the push for economies of scale has been driven by the Web's success, which has created demand for systems capable of handling thousands of transactions simultaneously with very low latency and sometimes, high data rates.

Open Internet and Web standards are the main enablers of the interoperability that allows very large independent data centers or "server farms" (typically operating 1,000's or more physical servers) to efficiently provide their vast processing power and storage capacity, via broadband Internet, to multiple customers.

"Virtualization" refers to the creation of virtual versions of operating systems, servers, storage and other computing resources that are hidden from the users. It supports the business models of data centers because it helps make it possible for customers to transparently offload their computing tasks and databases to those remote and massive computing facilities. Virtualization also enables data centers to internally reassign and reconfigure their processors and disks on the fly, making maximum use of those resources.

Consider elasticity, one of the essential characteristic of cloud computing. Many organizations have requirements to occasionally process massive amounts of data or specific projects have processing requirements that are beyond the organizations computer infrastructure to handle. In the past, the organization would either have to purchase more hardware and software or offload the processing to a contractor. Either solution can result in large expenditures and increases the risk of failure. In contrast, having access to a cloud computing platform that supports the organization's application means that when sudden processing requirements are needed, this can be offloaded to the cloud for processing. The net result is that the organization still has control of the processing, reduces risk, and reduces operational costs. Of course, there are now organizations that are moving entire applications to the cloud so that they do not need to deal with hardware procurements, hardware maintenance costs, and so forth.

Therefore, by providing virtual processing on the Internet, the cloud helps companies reduce their exposure to information technology risks, such as under-utilization, temporary surge of computation demands or early obsolescence of purchased hardware and software, and it helps them reduce operating costs such as system management, floor space, electrical power and cooling. The cloud also gives companies the option of shifting IT capital expenditures (CAPEX) to pay-as-you-go operational expenditures (OPEX), which offers financial flexibility with respect to paying for and writing off IT costs.

The cloud can provide environmental value as well as cost reductions. Today's very large data centers offer at least a 6:1 advantage (http://mvdirona.com/jrh/TalksAndPapers/JamesRH_Ladis2008.pdf) over small data centers in terms of economies of scale. In addition to financial savings, this translates into reduced carbon footprints and resource use, not only through reduced cooling and power demands, but also through an overall reduction in demand for computer hardware that has its own environmental footprint of embedded energy and mineral resources.

Gartner Research's October 2010 "2010 Hype Cycle Special Report" (http://www.gartner.com/it/page.jsp?id=1447613) reports that cloud computing overall appears to be just topping the peak in its hype cycle and private cloud computing is still rising. But they also note that the adoption and impact of cloud computing continues to expand, and they include cloud computing among the transformational technologies that will hit the mainstream in less than five years.

Growing demand for cloud-based geospatial applications and platforms

Cyberspace increasingly mediates our awareness of real space. Since the deployment of the first online mapping applications in 1996, Web-based map browsers, GPS enabled [2]applications, in-car navigation services, high resolution Earth imaging systems and mobile smart phone location applications have tremendously expanded the average person's awareness of "maps in computers". "Maps and apps" are very much a part of the smart phone market phenomenon. These technologies and products have also dramatically increased the volume of digital geographic (or "geospatial") data. In fact, the rate of capture of geospatial data is accelerating while the complexity of the technical and institutional arrangements that enable this data to be produced and made useful is increasing. Convergence of these technologies with each other and with the Web's base technologies creates a fertile platform for innovation, as we are seeing. The cloud plays an important role in innovation, because time to market for ideas is much faster when new companies don't need to invest time and money in providing basic computing infrastructure.

Convergence of technologies in the cloud and cloud computing are creating huge market pressures and can be viewed as truly disruptive. For example, "TaylorMade was able to adopt this software so quickly because it's not hosted on the servers at its headquarters in Carlsbad, California, but rather on remote computers in the cloud. It's a story that's happening over and over at many large corporations".

As with any IT community with elastic computing requirements or requirements to deploy applications at lower costs, the conventional geospatial markets can take advantage of the cloud. Geospatial processing often involves intensive computing and very large data sets, which explains why providers of geospatial processing resources and data resources are among cloud computing's early adopters. These providers apparent technical and market success indicates that other communities of interest using geospatial technology will almost certainly see widespread adoption of this new approach to information technology resource allocation.

There are a number of market drivers that are increasing demand for cloud computing. These drivers are just as important for the geospatial community as they are for the financial, retail, or other industries. A 2009 survey by Shavlik Technologies[3] and a 2010 Frost and Sullivan Report[4] found that:

- Data access and processing by all constituents;
- Server and licensing consolidation;
- Disaster recovery functionality;
- Operational elasticity;

[2] MIT Technology, Novermber 2010. http://www.technologyreview.com/business/26641/?nlid=3738
[3] http://virtualization.sys-con.com/node/1179528
[4] http://www.bradenton.com/2010/11/08/2719760/research-and-markets-cloud-computing.html

- Rapid scalability;
- Dramatically reduced IT costs associated with cloud computing.

are the leading drivers behind new investments in virtualization and cloud computing technology.

In the geospatial world, there are a number of applications for which cloud computing is well suited.

- Modeling: First, the complexity of geoprocessing models requires large computing capacities but on an intermittent basis. Examples of models heavily dependent on the use of multiple geospatial data sources are hydrology flow models, plume modeling, weather forecasting, and ocean current modeling.

- Fusion: There is an exponential growth in deployed Internet-connected sensors results in even greater exponential growth in location-referenced sensor data. At the same time, there are massive stores of "traditional" GIS data and other geospatially enabled resources, such as location enabled internet packets. There is a requirement to "fuse" or combine all of this sources and resources into new forms to improve situational awareness, decision making, and consumer experience.

- Enterprises and government ministries and agencies seek to create information through sophisticated mining based on geospatial criteria of the above-mentioned data stores and data streams, and this often requires extraordinary compute power, memory and storage.

- Scientists are beginning to use these new data sources in data-driven science. Scientists will use cloud computing for cyberinfrastructure-intensive virtual experiments, simulations, archiving (open data will drive growth in archiving), and networks of geolocated sensors. Cloud-delivered ease of exploring spatial relationships among Earth features and phenomena will likely lead to a significant increase in such virtual exploration activity. The economics of cloud computing will also give granting agencies an attractive alternative to funding purchases of computing equipment, which in many cases is not fully utilized.

- Demand: The highly variable (but large) number of users that require access to complex geoprocessing. A well known examples is indoor/outdoor navigation in a 3d environment.

In summary, from a geospatial processing perspective, a standards based cloud computing platform resolves critical issues concerning the performance of applications and models. These issues include availability of sufficient computational resources to solve given computational problems in a timely manner. Performing remote functionality in a cloud infrastructure means that the service (cloud) consumer is able to allocate as much resources as required (e.g. sufficient disk storage, network bandwidth or amount of CPUs) and therefore can expect a process to be performed according to specific Quality of Service (QoS) parameters (e.g. to be finished in a specific time period)[5]. The allocation of sufficient hardware resources in advance can either be done manually by the cloud consumer or realized through defined set of rules according to which the cloud

[5] http://www.geoinformatik2010.de/public/abstracts/foerster.pdf

infrastructure has to allocate automatically additional computational resources (e.g. in the case of high request rates of a web service).

Standards add value in the cloud environment

One of the big issues facing the continued growth of cloud computing are industry standards that will enable cross platform interoperability, consistent security mechanisms, and content sharing. This is why the Open Cloud Consortium[6] formed in 2010. AS their mission states:

> *The Open Cloud Consortium (OCC) is a member driven organization that supports the development of standards for cloud computing and frameworks for interoperating between clouds, develops testbeds for cloud computing, and supports reference implementations for cloud computing.*

The OGC mission and vision is very similar but with a focus on geospatial standards that enable geospatial content sharing, integration of geospatial services into a variety of infrastructures, and so on. From this perspective, a number of OGC standards are already cloud computing ready.

Some of the geospatial applications and platforms currently offered as cloud-based Web services are dependent on proprietary interfaces and encodings, but most of these also depend on open interfaces and encodings. Openness fosters innovation, expands markets, and creates new opportunities and efficiencies for both providers and users. This can be seen in the Web itself, which is based on open standards, and it can be seen in cloud-based geoprocessing. The success of the early geospatial cloud offerings validates the OGC's 16-year effort to develop a framework of open and freely available geospatial interface and encoding standards and related best practices.

OGC standards empower technology developers to make geospatial information and services accessible and useful with any application that needs to be geospatially enabled, whether or not these services are provisioned via the cloud. Standards provide layers of abstraction that hide implementation details. That is, they enable easy interoperability and "loose coupling" as opposed to one-to-one integration and "tight coupling." The existence of a well-used foundation of international industry standards means that geospatial technology providers can focus their efforts on applications instead of basic service infrastructure. Users benefit from increased product choices, including niche products developed to be used on top of platform products that offer open interfaces.

Profusion of standards-based services enables developers to efficiently turn workflow use cases into service chains. Services hosted at the Software as a Service (SaaS) level on the cloud can implement OGC standards, enabling other services to link to, and thus "bind" to them, and enabling them to bind to other services. GIS, Earth imaging systems, location service systems and others can be hosted at the Platform as a Service (PaaS) level on the cloud. Here, too, service-to-service binding takes place, but through application programming interfaces (APIs) that reside on top of more comprehensive software architectures.

[6] http://opencloudconsortium.org/

Inside and outside the geospatial domain, change involves risk, so most technology providers and enterprises are evaluating cloud computing carefully, moving applications to the cloud slowly and experimentally. One thing potential cloud customers worry about is the availability and viability of cloud vendors, so they want cloud vendors to provide standard interfaces and encodings that make solutions less dependent on a particular provider. Providers are, thus, naturally claiming commitment to open standards (http://itmanagement.earthweb.com/netsys/article.php/3830701/Cloud-Computing-Firms-Yes-to-Open-Standards.htm
Customers like the idea that cloud computing lends itself to rapid configuration and reconfiguration of value chains in a “cloud ecosystem” (http://www.appirio.com/ecosystem/). In this environment customers and providers can mix and match services of many kinds: SaaS, PaaS and Infrastructure-as-a-Service (IaaS). To accommodate multiple SaaS providers, PaaS and IaaS providers tend to offer open, industry standard ways to connect to their services. Major IT companies’ main cloud offerings may be monolithic, one-stop-shopping clouds for their customers, but all cloud providers rely on basic Internet and Web standards and also new virtualization standards; these standards are key elements of customers’ and providers’ competitive strategies. Providers understand that their customers want the freedom to avoid lock-in, that is, they want to be able to choose services and switch from one to another on short notice. This is one of the cloud’s big attractions, and customers want it to apply to geospatial applications.

Look to the value chain

In cloud computing, standards play a particularly important role because they enable links in value chains through interoperability and choice:

- Standards provide flexibility to do business with a new cloud provider without excessive effort or cost.
- Standards enable multiple cloud providers, including niche providers, to work together to deliver value-added solutions. When their systems encounter unexpected demand, for example, cloud providers can shift loads to their competitors.
- Standards enable cloud providers to more easily meet the varied needs of different customers.

Security, virtualization, and service level agreements are some of the areas where interoperability is essential to cloud value chains. Location is another, because some of the services bought and sold in complex and dynamic cloud service value chains require or deliver geospatial information. Location services can be simple or complex, from "where I am" in a tweet to "which satellite can scan this flood zone soonest and will the airport be accessible by road at 09 00hrs?" Designers of cloud services need to think about the standard interfaces, encodings and best practices that enable "where" information to pass between systems. Here “systems” means cloud service frameworks and also customer applications that support decision-making based on the location, motion and proximity of people, places, things and phenomena (such as temperature).

Information system architects should be aware that even the simplest location parameters can be passed in many different and incompatible ways. Simple latitude/longitude location, for example, can be passed in RSS (Really Simple Syndication) messages in at least three ways. Though most GPS-based navigation and location services now use spatial reference systems based on the “WGS-84” mathematical model of our not-perfectly spherical Earth, there are important datasets that use other models, and there are literally thousands of spatial reference systems in use around the world. Interoperability is made more difficult by the fact that there are multiple technologies for representing location information (raster, vector, triangulated irregular networks, point clouds, computer-aided design Euclidean geometry and others). Semantics are a big issue, because different user communities define and name spatial features differently ("Is it a cul-de-sac or a dead-end?"). And there are many different proprietary interfaces, encodings and formats used in vendors' GIS (geographic information systems), location services, Earth imaging, facilities management and navigation products.

Whether simple or complex, if services can't exchange location information, the cloud service value chain breaks. (This is mostly a problem for SaaS and PaaS, and not for IaaS, because infrastructure resides at an architectural level below most geospatial communications.) Communication among software and data services usually requires that applications share common interfaces and encodings, and sometimes this requirement extends to the software platform that supports the applications. The geospatial interface and encoding standards developed in the Open Geospatial Consortium (OGC), often in coordination with other standards organizations, provide these open interfaces and encodings.

Is it GIS?

Prior to the wide use of OGC's geospatial interface and encoding standards, most geospatial information was confined within Geographic Information Systems (GIS) and Earth imaging systems and their specialized spatial databases and the networks of users who used software from the same vendor. Open OGC standards first enabled different vendors' GISs to exchange instructions and data. Then they enabled database software vendors to accommodate all types of spatial data in their products, and they enabled improved "fusion" of vector and raster spatial data types.

The geospatial standards framework has continued to evolve and expand to support mobile communications, map browsers and sensor webs. (Every sensor – including video devices, cell phones, wearable medical monitors, vehicle pollution controls, RFID chips, etc. – has a location, and location often matters). As a result, though there is still a growing need for the kinds of complex spatial processing conceived, packaged and marketed as GIS, there is an exploding need for spatial data and simple kinds of spatial services, such as "Get a map," "Where am I?", "Where is the nearest pizza shop?", and "Is it safe to dig here?" These are now typically provided without a GIS. The geospatial industry – previously referred to universally as the GIS industry – is undergoing rapid change because a growing number of user communities need a growing number of geospatial services, and all such services, from simplest to most complex, can now potentially be offered as Web services that plug into otherwise non-spatial applications.

These Web services can now be conceived, packaged, marketed – and sometimes most efficiently delivered – as cloud services. OGC Web Services (OWS) standards provide the common "glue" that makes geospatial Web services practical as geospatial cloud services.

In the new context of Web services, there is a much freer flow of location information between domains. Previously this information exchange depended on batch conversion of files between different spatial formats. GIS practitioners are getting used to the new reality of using “services” instead of “files”, and file management expertise is becoming less important in their work. Even the term geospatial is yielding to the more general terms spatial and location, as efforts to remove the technical barriers between indoor and outdoor location gather momentum. Spatial data has become just another data type, useful in all kinds of applications.

Spatial Data as a Service (DaaS)?

Spatial data is usually expensive to collect and manage, and its value usually increases with the number of people who can use it. The cloud reduces the cost of data distribution, through what is sometimes described as “Data as a Service (DaaS)”. This offers governments and businesses a way to increase the return on their expensive spatial data investments by making the data available for more applications.

Technically, however, the term DaaS is misleading. In a review of this paper, Doug Nebert, chief architect for the US Federal Geographic Data Committee’s GeoCloud Sandbox Initiative (see below), explained, “DaaS is a recent invention that is not resident in the cloud IaaS/PaaS/SaaS service stack. It does not have a clear definition, other than being just another application providing data access services, which our (GeoCloud Sandbox) PaaS offerings do. It is inaccurate to describe Google and Bing as offering DaaS, since they are not providing access to data, but rather visualizations of data. They allow users to serve and integrate their own KML files (through proprietary APIs), but only a tiny minority of information is actually served by these giants as data in KML.”

Keeping in mind that DaaS is actually SaaS that delivers data, it is nevertheless true that sources of data, free or otherwise, can become more useful when they are accessible through Web services and provisioned on the cloud for reliability and performance. Google Public Data Explorer (http://www.google.com/publicdata/home) and Wolfram Alpha (http://www.wolframalpha.com/), for example, provide ways to utilize a wide range of free Web-accessible datasets in various ways, including map displays and limited spatial analysis.

Cloud-based services can help businesses unlock the value of their data assets, utilizing location intelligence applications to derive insight that improves competitiveness and business performance. SaaS applications can make corporate data available to more applications and more users within the company and make corporate data available in the same user environment as other kinds of data from other sources.

The FGDC GeoCloud Initiative

Similarly, governments want their investments in data to provide the greatest possible value to the taxpayers. Cloud computing efficiently leverages the investments that government agencies have made in OGC standards development and service oriented spatial data infrastructure (SDI) architecture development over the last decade.

The US Federal Geographic Data Committee’s GeoCloud Sandbox Initiative is a one-year initiative to explore the hosting of government geospatial data and geospatial services in vendor-provided and agency clouds. The FGDC’s goal is to identify requirements-driven solution architectures for geospatial data and services and to document scalability, reliability, costs and redundancy. Data services will be provided first, then processing services, and then solutions as part of realizing a Geospatial Platform environment that coordinates access to geospatial assets in a broad portfolio across government.

The GeoCloud Sandbox initiative is part of a larger initiative in the US Government to modernize government IT strategies. The GeoCloud Sandbox will use the US General Services Administration's apps.gov IaaS services that provide what is essentially cloud-based "shared hosting" of networks, raw computers, and associated operating systems. On top of that infrastructure, the GeoCloud Sandbox intends to pilot the deployment of candidate services or solutions architectures (suites of software) to provide common geospatial capabilities for government agencies. Another objective is to monitor the real costs of operational hosting for different configurations to inform agencies about appropriate options for data and services deployment. The expectation is that agencies will be able to reduce the cost and complexity of buying and managing hardware and software.

Spatial data infrastructure initiatives in other parts of the world will likely begin similar initiatives to investigate the value of moving government geospatial resources to the cloud. INSPIRE (INfrastructure for SPatial InfoRmation in Europe), for example, requires guaranteed response times for specific queries, and the cloud provides scalable and affordable solutions capable of meeting such requirements.[7] The economics of information technology management compel government data managers to consider the IT performance improvements and management time and cost savings that are potentially available with cloud computing.

One example of this in Europe is the European Space Agency (ESA) G-POD Cloud project[8]. Initiated in July 2010, it aims at validating the idea that cloud computing can be used efficiently to process large amounts of earth observation data. Taking advantage of the inherent transportability of ESA Grid (G-POD) processing jobs, it coupling G-POD with cloud resources to perform complex EO data processing jobs in a sustainable and cost-effective way. While the G-POD infrastructure uses mainly Grid protocols, models and resources for processing and data access, the basic commitment to application encapsulation and virtualization has made porting to the cloud relatively easy.

Reasons for caution

Businesses and agencies, large and small, may see the cloud as the future, but most are entering into this domain one step at a time, assessing possible problems, as well as opportunities.

Pricing and service level agreements are one issue. Is there a good match between your actual usage patterns and the cloud services that are offered? If a cloud provider, for example, is billing by the full hour for pay-as-you-use, and you're using the service for only a few minutes out of the hour, your use may not justify the cost.

Perhaps the most commonly cited concern is security. Cloud service providers may offer better security at a lower cost than many small and medium size businesses might be able to contract otherwise. Nevertheless, customers should learn about the security of

7 "Towards Spatial Data Infrastructures in the Clouds" Bastian Schäffer, Bastian Baranski, Theodor Foerster

8 csse.usc.edu/gsaw/gsaw2010/s11d/brito.pdf

prospective data centers and ask whether cloud service providers conduct background checks of employees and provide adequate security training and IT auditing. Cloud vendors may use third parties to host data centers and hardware, and this increases security risks. They should also be aware of the terms and conditions – what are they signed up for?

What's needed and what's next?

The cloud is based on a standards framework for service oriented architectures that provides for "publish, find, bind":

1. Publish: Resources can be hosted and their description, network location and interfaces can be published in standards-based registries or catalogs.

2. Find: Client applications can search the registries or catalogs to find a resource.

3. Bind: The client application can invoke the server through standard interfaces.

What's missing in this scheme is "Agree." The "Publish" activity needs to be able to provide metadata describing, for a particular service or data resource, details about authentication, authorization, confidentiality, integrity, non-repudiation, protection and privacy. Servers need to be equipped to manage these issues, and owners of data and services need the tools and understanding necessary to configure controls.

The OGC has done some of the technical standards work that provides the foundation for "agreement management" that will bring us closer to full realization of geospatial cloud applications:

- The Geo Rights Management Domain Working Group (GeoRM DWG) (http://www.opengeospatial.org/projects/groups/geormwg) has produced the Geospatial Digital Rights Management Reference Model (Abstract Specification Topic 18). The mission of the GeoRM Working Group is to coordinate and mature the development and validation of work being done on digital rights management for the geospatial community. The reference model has been approved by the OGC membership, who will use the GeoDRM RM in developing OGC standards for open interfaces and encodings that will enable diverse systems to participate in transactions involving geospatial data, services and intellectual property protection.

- The Security DWG (http://www.opengeospatial.org/projects/groups/securitywg) is a forum for discussing topics related to authentication, access control and secure communication.

- The OGC GeoXACML Standards Working Group (http://www.opengeospatial.org/projects/groups/geoxacmlswg) developed GeoXACML, which is based on the OASIS XACML standard. XACML (eXtensible Access Control Markup Language) was developed by OASIS (the Organization for the Advancement of Structured Information Standards).

- The Workflow DWG (http://www.opengeospatial.org/projects/groups/workflowdwg) addresses geospatial workflow issues, including security and licensing issues such as data encryption, authentication, and provenance tracking. Members of the Workflow Management Coalition http://www.wfmc.org/, an OGC Alliance Partner, participate in the OGC Workflow DWG as part of their effort to foster open standardized workflow methods.

Also, in 2009 the OGC Board of Directors created the OGC Spatial Law and Policy Committee (http://www.opengeospatial.org/pressroom/pressreleases/964) to provide an open forum for OGC members' legal and policy advisors to discuss the unique and increasingly critical legal and policy issues associated with spatial data and technology.

Conclusion

OGC Web Services (OWS) standards were developed to make geospatial data and services an integral part of Web-based distributed computing, and so these standards are ready-made for cloud computing. They define the open interfaces and encodings that are needed to successfully host all types of geoprocessing on the cloud. This provides a variety of opportunities for large and small enterprises, government agencies focused on building spatial data infrastructures, and traditional vendors of geospatial software and data. The efficiencies afforded by the cloud computing model will contribute to growth in most geospatial market sectors, including mass market consumer applications, sensor webs, location services, Earth observation data markets, CAD/geospatial 3D modeling, and municipal management.

Though many data and service providers are making a move into cloud computing, most adopters are moving cautiously, testing the various advantages over previous provisioning models and looking for assurances regarding licensing and security of data hosted by cloud providers.

Much remains to be done. Science, commerce, government, education and everyday casual users cannot benefit fully from today's technical information infrastructure until the social, institutional, behavioral and commercial parts of the information infrastructure have matured. The human aspects of the infrastructure will become easier to deploy as solutions and tools become available that implement emerging interface and encoding standards for service chaining, data licensing, rights management, pricing and ordering, and order fulfillment. If participants in related standards efforts work from a complete set of requirements, and if they coordinate their efforts to avoid standards gaps, redundancies and inconsistencies, the convergence of technologies involved in cloud computing and geospatial data management will likely yield numerous benefits for society.