OSBR.CA

The Open Source Business Resource

MARCH 2009

MARCH 2009

PUBLISHER:
The Open Source Business Resource is a monthly publication of the Talent First Network. Archives are available at the website: http://www.osbr.ca

EDITOR:
Dru Lavigne
dru@osbr.ca

ISSN:
1913-6102

The editorial theme for the March issue of the OSBR is "Geospatial" and the role open source is playing in transforming this niche market into a mainstream market. The authors in this issue bring their many years of experience in both industry and open source to provide their observations, lessons learned, and to provide examples of open source geospatial implementations. Even if you don't use geospatial technologies, you'll still find many valuable insights in this issue of the OSBR.

As always, we encourage readers to share articles of interest with their colleagues, and to provide their comments either online or directly to the authors. We hope you enjoy this issue of the OSBR.

The editorial theme for the upcoming April issue of the OSBR is "Open APIs" and the guest editor will be Michael Weiss from Carleton University.

Dru Lavigne

Editor-in-Chief

dru@osbr.ca

Dru Lavigne is a technical writer and IT consultant who has been active with open source communities since the mid-1990s. She writes regularly for O'Reilly anddNSStuff.com and is the author of the books BSD Hacks and The Best of FreeBSD Basics.

The recent emergence of MapQuest, Google Earth, Garmin GPS, and many other modern geospatial products make it seem that mapping technologies are a relatively new component of today's information technologies. In fact, the mapping industry was one of the original adopters of technology when geographic information systems were first developed over 40 years ago. The fruits of this backroom technology, once the domain of highly trained specialists, is now being leveraged by hundreds of millions, if not billions, of consumers around the world.

Open source geospatial technologies have followed this same path from niche technology to mainstream component and are now critical to many of the applications that business and consumers use on a daily basis. Google Earth, for instance, incorporates a critical component of the open source geospatial stack to deliver satellite imagery to several hundred million installations around the planet. In the following articles, you will learn more about how these critical niche technologies have evolved from small grassroots activities to thriving technology projects under the umbrella of the Open Source Geospatial Foundation, and ultimately into key commercial components of industry.

Paul Ramsey, Senior Consultant with OpenGeo provides us with an overview of the state of open source geospatial technologies within the geospatial industry. Paul explores the strengths and weaknesses of today's open source geospatial stack and provides an indepth background into how these technologies have evolved to their current state.

Tyler Mitchell, Executive Director of OSGeo provides readers with an overview of the critical role OSGeo is now playing as an enabler of community and technology growth. Tyler discusses the many technical, business, and perception barriers that OSGeo is successfully lowering for broader adoption of open source geospatial technologies.

Mark Lucas, Principal Scientist, and Scott Bortman, Senior Software Engineer at Radiant Blue Technologies, provide us with an overview of OMAR as a collection of open source geospatial technologies that are experiencing broad adoption within the US Department of Defense.

Andrew Ross, Director of Engineering at Ingres Canada, reviews the process by which Ingres recognized the need for geospatial capabilities within the Ingres database. He also discusses collaboration with the OSGeo community to meet Ingres' technology needs while at the same time contributing back to the community.

Haris Kurtagic, General Manager of SL-King, and Geoff Zeiss, Director of Technology at Autodesk, discuss their experience in implementing RESTful services with geospatial technologies, and the important role standards are playing in the development of geospatial web services.

Noticeable among these articles is the diversity of perspectives and uses for which geospatial technologies are being adopted today. Interestingly, the geospatial industry is struggling to define itself as information technology begins to integrate geospatial capabilities into stacks of technologies, pushing the limits of the traditional geographic information systems (GIS) industry. OSGeo and the open source geospatial community are unique in that open source has become the defining aspect and strength of this community as it grows at rapid pace into a dominant force for the provision of geospatial technology within industry.

Dave McIlhagga is the president and founder of DM Solutions Group Inc., a leader in web mapping solutions delivery since 1998. Dave has positioned DM Solutions Group as a leading provider of commercial products and services to the open source web mapping community. Recently he has led the company's effort to bring high quality custom mapping to consumers through MapSherpa.com - to be launched in spring of 2009. Dave is a former Board member of the Open Source GeoSpatial Foundation and an active contributor to the open source geospatial movement. Prior to founding DM Solutions Group, Dave was a leading developer of one of the industry's first web mapping technologies at TYDAC Research. Dave graduated from Carleton University with an Honours Bachelor's degree in Geography, concentrating in Geographic Information Processing.

"Man is a microcosm, or a little world, because he is an extract from all the stars and planets of the whole firmament, from the earth and the elements; and so he is their quintessence."

Philipus Aureolus Paracelsus

Open source has seen great success in general information processing, but does it have a future in vertical markets? In this article, we examine how geospatial open source provides an example of the market challenges of a mid-sized vertical market.

Open Source and Decision Makers

General purpose open source software (OSS) has been on the radar of decision makers for almost a decade. Big projects like Linux, Apache, Firefox and Open Office are supported by Fortune 500 companies like IBM and Sun. Everyone knows about open source, it is in their plans, books have been written. In general information technology (IT), there is little more to say about open source.

However, the IT world does not end at databases, operating systems, and office automation. IT is like a fractal form, each major facet can be subdivided and re-subdivided down into particular shapes that fit the needs of unique markets.

The economy is full of niche markets with very particular information processing requirements. Examples of niche markets can be found in health care, education, natural resources, manufacturing, and telephony. Each of these fields makes use of the generic open source building blocks that have already swept through IT. They also have their own distinct ecosystems of dominant proprietary vendors and de facto standards that shape decision making and software acquisition.

Open source is growing in these niche markets, but much more slowly than in general purpose IT. The reasons are pretty straightforward: smaller markets mean fewer users, fewer developers, and fewer resources for open source.

History of Geospatial Markets

Geospatial is one of the niche fields that is being slowly colonized by open source. Geospatial software is used by natural resource managers, cartographers, fleet managers, and anyone with a location component in their data.

Geospatial is a recent term for what has traditionally been known as "geographic information systems", or GIS. The original users of GIS software were government environmental scientists and land managers. As early as the mid-1960s, governments were building their own GIS systems, writing the code in-house (http://www3.nfb.ca/collection/films/fiche/?id=18208).

In the 1980s, as computing hardware became cheaper and more interoperable, the economics of GIS shifted. Rather than buying computers and writing GIS systems from scratch, governments bought computers and then bought GIS software. An ecosystem of proprietary GIS vendors emerged quickly. Some vendors were regionally based, others were specialized in particular fields like forestry or oil and gas.

Some of the last in-house systems were written by the government. The Geographic Resources Analysis Support System (GRASS, http://en.wikipedia.org/wiki/GRASS_GIS) was written by the US Army Corps of Engineers after an evaluation process that determined none of the current proprietary systems met their needs. The Map Overlay and Statistical System (MOSS, http://en.wikipedia.org/wiki/Map_Overlay_and_Statistical_System) was written by US Fish and Wildlife Service, after a similar market evaluation.

Both GRASS and MOSS were released as public domain works, effectively becoming the earliest examples of open source geospatial software. GRASS was re-licensed as a GPL project in the mid-1990s. However, throughout the 1990s, the OSS largely languished in academia while proprietary software filled the entire government and corporate ecosystem for workstation-based GIS.

Just as the computer and operating system market consolidated to Intel and Microsoft through the 1980s, the GIS market consolidated through the 1990s. The GIS market consolidation battle in North America and, increasingly, the rest of the world, was won by a proprietary vendor, ESRI (http://www.esri.com/), which started out from a market base in the US federal government and gradually displaced other competitors in the North American market. By 2000, ESRI had achieved a geospatial market position similar to that of Oracle or Microsoft in the general IT marketplace. They were the default vendor and the only safe choice for decision makers.

Like Microsoft and Oracle, ESRI's market dominance was built on a narrow but important product category. Microsoft dominated operating systems and office automation. Oracle dominated relational databases. ESRI's dominant category was, and still is, desktop GIS software. Desktop GIS software provides users with the capability to create, edit, analyze, and cartographically print geospatial data.

Market Disruption

The rise of the Internet was famously disruptive to the Microsoft business model. In 1995, Bill Gates radically revised the company's software strategy to focus on networked communication (http://www.usdoj.gov/atr/cases/exhibits/20.pdf).

The new strategy was intended to preempt new competitors arising who could take advantage of the Internet software arena. On the desktop side, Microsoft was successful, Netscape was crushed, and Real Networks is mostly gone.

In geospatial, ESRI suffered a similar disruption from the Internet. As their users got used to accessing non-spatial data over the Internet, they began to ask an obvious question: "how can we provide access to our geospatial data over the Internet?".

The immediate question didn't involve analysis, cartographic printing, or data capture, all of which are ESRI's core desktop strengths. The immediate request was for simple data access. Like Oracle, ESRI made some early strategic errors in providing Internet-enabled versions of their software. Their pricing model locked smaller organizations out as the cost of entry was too high. However, rather than throwing up their hands and not providing Internet services, smaller organizations and individuals simply looked for alternatives. In the case of Oracle, they found MySQL (http://www.mysql.org) and PostgreSQL (http://www.postgresql.org). In the case of ESRI, they found MapServer (http://mapserver.org).

Arguably, the price points that individuals and small organizations wanted couldn't be rationally provided by ESRI or Oracle. But, by driving individuals and smaller organizations away, the dominant vendors seeded a new marketplace that quickly developed alternative open source product suites.

Dynamics of Open Source Geospatial

In 2000, the market for "web map servers" was brand new. The high cost and poor performance of the offerings from ESRI slowed adoption.

The low technical bar to entry needed to "display maps on the Internet" allowed a new entrant to gain traction. MapServer was originally developed in an academic setting, the University of Minnesota, using grant money. Released as open source in 2000, it had just enough existing functionality to begin attracting new developers.

Once MapServer began providing useful services to users, it started to attract more open source development. Unlike the "pay to play" model of proprietary products, the open source MapServer allowed organizations to get their feet wet with existing capabilities for free, then pay to add new capabilities if and when they needed them. Organizations that adopted MapServer for free began using funds to add improvements around the margins.

Among the features that were added by the MapServer user community were the ability to read data from additional GIS file formats and from satellite imagery, and support for international interoperability standards. The improvements, and particularly the standards support, served to make MapServer useful to a larger audience, which drove market growth even more.

The characteristics of the adopters of open source geospatial are familiar to any student of open source market dynamics. Open source tools are generally evangelized by technical staff, who have the ability to acquire and test the tools themselves. This effectively limits early adoption to startups and larger organizations with pockets of progressive technical expertise. Conservative organizations tend to gather intelligence through vendors and trade magazines, which serve mainly to reinforce existing purchasing patterns.

Organizations that have already chosen a proprietary product for a functional category will rarely switch to an open source equivalent until the proprietary product hits end-of-life.

The exceptions to this rule have been open source phenomenons, like Linux or Firefox, which have received publicity outside the technical trade press. The publicity around Linux has forced even conservative organizations to make some formal consideration of using general open source.

Open source geospatial is a small niche and its software will never receive the popular press coverage of the sort that has made Linux and Firefox well-known. Most existing geospatial software customers, such as governments and resource industries, remain comfortably within the arms of the dominant proprietary vendor, ESRI.

The market area where open source geospatial has been most quickly taken up is among organizations in which there is no existing bias towards the dominant proprietary vendor. These are usually companies or agencies making their first foray into Internet mapping, motivated by the geospatial renaissance triggered by the introduction of Google Earth and Google Maps. Their technical staff research the options and find that they can get what they want from open source. The available open source includes:

- the PostGIS (http://www.postgis.org/) geospatial database
- GeoServer (http://www.geoserver.org/) to provide web map services and web feature services
- the GDAL (http://gdal.org) programming library

- the Geotools Java GIS toolkit (http://geotools.codehaus.org/)

- the SharpMap (http://www.codeplex.com/SharpMap) geospatial application framework

- the web interfaces provided by Open Layers (http://openlayers.org/)

Providers of web-based services, which have to scale their infrastructure based on potentially exponential changes in users, are particularly enthusiastic about open source. The capital cost of scaling an infrastructure that uses proprietary software is $X dollars times the number of nodes. For open source, the capital cost of scaling is zero. The annual infrastructure support costs for proprietary software are again locked linearly to the number of nodes. For open source, the support costs in consulting fees to open source programmers or companies generally increase with the number of nodes, because the odds of needing new features or exposing bugs increase as use increases, but the increase is generally less than linearly.

Examples of new Internet companies using open source geospatial include:

- Redfin (http://www.redfin.com/), a real estate information company

- GeoCommons (http://www.geocommons.com/), a data sharing community

- Zonar (http://zonarsystems.com/), a fleet management and vehicle tracking company

- GlobeXplorer (http://www.globexplorer.com), a satellite imagery re-seller

- Urban Spoon (http://www.urbanspoon.com), a restaurant review site

- Whereyougonnabe (http://www.whereyougonnabe.com/), a spatial add-on for Facebook

The other area where open source geospatial has been taken up is in technically savvy pockets of large organizations. Like the Internet companies, these organizations have enough of a user-side demand that deploying proprietary Internet service software creates a large financial burden. Unlike startup companies, large organizations can potentially afford to pay for the proprietary software.

In some large, conservative organizations, visionary managers adopt and deploy open source. However, there is nothing systematic about the adoption, it is a consequence of the personality and personal interests of the manager. If he or she leaves the company, the open source pocket may disappear. Because the progress of open source geospatial in these organizations is so personality-based, it tends to be rare and to run in fits and starts.

Open Source Geospatial Challenges

The progress of open source geospatial on the desktop has been very slow. Desktop geospatial software already has an entrenched proprietary incumbent, ESRI's ArcGIS (http://www.esri.com/software/arcgis/), with a long history. The amount of quality code required to reach feature-parity with the incumbent is very high as ESRI has been working on their desktop software for decades.

Simple desktop implementations are available with QGIS (http://qgis.org), uDig (http://udig.refractions.net/) and gvSIG (http://www.gvsig.gva.es/), but are mostly consigned to the niche of extremely low budget organizations.

As a result, financial resources are not available to speed up development, and the pace of progress remains slow. The exception has been gvSIG, which is heavily funded by Spanish government organizations, but it is still mostly a niche development used in Spain.

In all cases, the growth of open source geospatial has been slowed by matters of scale. Open source products generate much smaller revenue streams from user populations than proprietary products. In large markets with well-capitalized customers, companies can make good money even on the smaller revenue flow of open source.

However, in a small vertical market, it is difficult for companies to get a foothold. A customer will usually deploy several open source geospatial products to create a solution, so a support provider has to have extensive in-house experience to support the whole solution. In a traditional model, start-up costs would be capitalized by a venture funder, but the size of the geospatial market-place makes the 10:1 returns required by venture capital unlikely.

The OpenGeo Foundation (http://opengeo.org) is breaking the geospatial support log-jam by building a social enterprise using philanthropic funding to bootstrap an organization that contains the breadth of expertise necessary to support a variety of open source geospatial applications. OpenGeo's motivation is not to maximize profit, but to maximize social good, while covering costs. This allows the organization to build a sustainable market while surviving on the smaller revenue streams available in the open source geospatial arena. The products OpenGeo supports such as Geoserver, OpenLayers, PostGIS, and GeoExt (http://www.geoext.org/), make a top-to-bottom deployment stack for geospatial applications.

Lessons for Other Markets

Open source geospatial holds a number of lessons for other vertical markets. First, frontal assaults on the leading proprietary vendor are unlikely to succeed. In their core areas, the leading vendor has an advantage in technology development and existing mind-share. Usually, building enough technology to compete with a leading vendor head-to-head takes years of development, and a partially functional product will be ignored.

Second, disruptive changes in technology provide opportunities for open source. Most leading vendors carved out their advantage on the desktop during the 1980s and 1990s. The transition to web-based services has opened a temporary gap in the marketplace where existing vendors have a smaller technology advantage, and their marketing advantage is limited to their existing universe of customers. Open source can become the core for new service-based companies competing with proprietary software vendors.

Finally, new markets for capabilities are the most fertile opportunity of all. In geospatial, its expansion into daily life, through vehicle and device tracking, low cost aerial imaging, and handheld mapping, is growing the market exponentially. New developers and managers, without long-held preconceptions, are making technology choices. On a level playing field, open source Internet technologies are regularly winning.

Paul Ramsey is a Senior Consultant with OpenGeo, a social enterprise dedicated to providing support and service on a complete stack of open source geospatial products, from database to middleware to web interface. Mr. Ramsey is a Director of the Open Source Geospatial Foundation and the 2008 recipient of the Sol Katz Award for accomplishments in open source geospatial software.

"The logic is compelling; depending on closed source code is an unacceptable strategic business risk. So much so that I believe it will not be very long until closed-source single-vendor acquisitions when there is an open-source alternative available will be viewed as actual fiduciary irresponsibility, and rightly grounds for a shareholder lawsuit."

Eric S. Raymond
http://www.catb.org/~esr/writings/cathedral-bazaar/magic-cauldron

Furthering the adoption of open source software (OSS) is often seen as a natural, contagious progression as more developers and users share with others about the success of their projects. But how can an open source project even hope to compete with proprietary commercial products with massive marketing budgets and staff? Aside from not typically having large financial resources for marketing, other factors can lead enterprise users to look elsewhere for guarantees of product longevity and a robust support ecosystem around the product. Without these features, many users and businesses alike would not consider the software as an option or have a desire to become involved.

This is equally true with open source geospatial software projects which are focused on presenting freely available mapping and geographic analysis tools to the world. The Open Source Geospatial Foundation (OSGeo, http://osgeo.org) seeks to address the needs mentioned above, to promote the excellent software that is available, and to provide a model where businesses can join in promotion and development. OSGeo undertook some novel approaches to encouraging new and existing support options in order to boost confidence within the business sector and ensure that project code will be publicly accessible for years to come.

In this article, we discuss the factors needed to get open source geospatial products into the hands of those users willing to test, use and eventually admire them as their favourites, or to go one step further and recommend them as a corporate solution. We examine the advantages provided by using marketing to help promote open source projects and then consider how this can boost business confidence in the use of the software. The article closes with a look at how a natural ecosystem of open source users is able to create something bigger and more consequential than each project could attain on its own.

A Herald of Success

For more than a decade, users of geospatial software have been able to turn to various open source products to help them do their work. In particular, projects such as the GRASS (http://grass.osgeo.org/) geographical information system (GIS) workstation product have aided analysts and researchers in the advanced analysis of geography. As the web emerged, projects such as MapServer (http://mapserver.org) materialized and helped analysts share their maps with the world. To this day, new and innovative open source tools continue to address the needs of analysts in a productive manner. As a result, there is a growing desire to share these advances and success stories with others who have yet to hear them.

Why does it matter that users want to see others share in their interest in a project? Often, like attracts like, and one person may share with another who would appreciate such software. As these new users become involved, they may help to identify bugs in the software, contribute to documentation of the project, provide user support, or even become contributing developers.

Some will go on to develop their own consultancy business or bring tools into their existing organization.

The open source model provides several ways for keen, new users to help make the project better and many projects actively seek new involvement. In salesman talk, new users are easy customers to sell to as no cold call is required. New users encourage others. The end result is a stronger project, with an increasingly large user base and a higher profile image which attracts yet more users and builds an even stronger user base. If a project is to remain viable in the foreseeable future, it depends greatly on its ability to propagate new users and settle into new organizations.

There also exists those scenarios where a stranger using another product may never realise an alternative exists, even one that may prove more effective. The typical consumer may never communicate back to the original project team but instead silently enjoy the fruits of others' labour. For many projects, having user success stories made public to other potential users is critical to advancing its user community and showing a strong footing in the larger software landscape. This helps to inform people who might otherwise never learn about the product that has been freely and openly available to them. This open sharing of experience not only helps existing users, but encourages new ones to look further into it.

Marketing Open Source

Marketing can be a powerful tool. The proprietary sector knows this and relies on it to sell their products to those willing to pay. As large open source products emerged and developed around new or existing commercial companies, open source marketing also started to develop.

With large and respectable organizations helping to support the marketing efforts of an open source product, more users hear about the product and the risk to investigate or use the product within their organization is reduced.

Open source products that are tied to particular commercial entities are both blessed and cursed. They are blessed that they have easier access to capital funding for marketing and cursed that they less easily attract new project team members because volunteers tend to give to causes that are less commercial.

These are only generalizations, but it is common sense that the open source projects that have a marketing budget or stable corporate entity behind them can market the project in many sophisticated ways. This is especially the case when the potential client is not a casual user, but rather a professional representing a larger organization with whom corporate marketing gets more traction.

The message of marketing in open source can be slightly different than in the proprietary world. For example, some campaigns may focus on the aspect of having the software be available at no cost. But in the end, adopters of a product need to know two things: that the product will run successfully in the future, and that support is available on a number of fronts. Marketing plays into both of these factors, but does not merely advertise them. The administrative and organizational strength behind the projects must be exhibited to assure longevity and to promote a healthy support ecosystem. If these are in place, projects have a much higher chance of gaining widespread enterprise usage, attracting more experienced developers, and improving their prospects for the future. Demonstrating these strengths is part of the marketing process.

This approach treats the open source project as a product, one that needs to be sold to end users. While in traditional product sales the end result is a fee paid for a license, in OSS there is no financial transaction to serve as a direct commitment. Thousands of users may download and use the product on a regular basis, but there is no need to register or self identify as users. In this case, the sale was successful, but there is no clear metric to account for it. Even so, potential end users still need to learn that a product exists and that it is ready to meet their challenges.

These requirements are not unique to open source. Consider the case of a new bank. It may spend millions on marketing to potential clients, but if the foundation of the building is not laid, the customers have nowhere to go. Likewise, simply setting up a tent in a local park would fail to build the long-term confidence that is necessary in banking. Furthermore, if a potential customer would like to set up an account and the bank has no staff, then confidence would be diminished.

OSGeo

Addressing the organizational and marketing needs of several open source projects was a key goal when starting the OSGeo Foundation. The organization itself was created to help assure project longevity, encourage a healthy support and user ecosystem, and act as a focal point for various communities to come together for advancing common goals.

OSGeo was started in 2006 as a nonprofit organization, interest having grown over a year or two prior on several fronts. There had been a general recognition that various software projects were very mature and used as stable solutions to geographic planning and mapping exercises.

While competition with proprietary products was very real, the uptake of the open source solutions was steadily continuing. The question in several minds was, "How can we further advance these great products we use?" It was assumed that, eventually, members from the user and consulting community might find some way to gain exposure to some projects. What no one could expect was that a large corporate entity had been working toward releasing their proprietary product as open source. This development was the fuel to move the vision of a formal organization forward.

Autodesk (http://autodesk.com), known for their Autocad design and media applications, also produces geographic information management products, including a popular web-based mapping tool, MapGuide (http://mapguide.org). MapGuide was the first Autodesk product to be released as open source.

Being already familiar with existing open source geospatial projects and the community development approach, Autodesk sought to work cooperatively with these projects rather than release their product as a competitor. Several other projects had already been available for a while, but only a few had any corporate presence behind them. MapServer had been hosted by the University of Minnesota, but had extensive contributions from external companies. This meant that Autodesk was able to find companies who were already ardent developers of parts of MapServer. The existence of companies that support particular open source products helps other businesses to have confidence to investigate them further.

In finding a good way to work together, it was proposed to develop a non-profit organization to help focus on common needs and goals across many projects.

All projects need technology infrastructure such as web servers, code repositories, and mailing lists. All could benefit from collaborative marketing in venues where it might not be feasible for a single project to go alone. These and other concepts brought two dozen leaders of, and contributors to, open source geospatial software together for a meeting. The outcome became OSGeo, with a board of directors well known in the industry, dozens of charter members, nine specific software projects dedicated to working together, and Autodesk as the founding sponsor. Three years later, with over 70 charter members, dozens of local chapters spread around the world, and over a dozen sponsors, OSGeo is addressing significant issues.

A Gateway For Business

There were three specific ways that the development of OSGeo lead to increased business confidence in its open source product offerings. These were:

- heightened confidence in embracing the software

- greater certainty that code contributions would be well invested

- a way to contribute financially to the open source products

There was general agreement that all projects joining OSGeo would find some benefit through increased promotion and lowered costs for their technology infrastructures. It was also well known that end-users would have increased confidence in using a product with a formal organization behind it, as opposed to an ad-hoc project management structure with, ultimately, limited accountability. This role of OSGeo is perhaps the most significant as open source tools are not always weighed solely on their functional merits. Other factors include:

- availability of support, particularly documentation

- perceived total cost of ownership

- long-term viability of the tools

Often, potential users are limited by their ability to adopt a product that does not appear to be "serious" to their upper management or decision-makers. Having a formal organization standing behind the projects provides a reinforced perspective for those to whom a trusted name is important. They can have confidence that the software project management team is not a fly-by-night operation and that their product will be available in the future. Having such confidence is paramount.

As OSGeo stands behind its software projects, it encourages further development and support, both voluntary and commercial. It is hard to imagine anyone knowingly investing programming effort in a product destined to dissolve after several months or years. System integrators and consultants who use or contribute to an open source project look for projects that will remain freely accessible and hopefully flourish. Those who contribute to open source projects do not want to see their contribution get lost or locked up in a product that they cannot use or contribute to in the future. Requiring an OSI recognized license (http://opensource.org/docs/osd), requiring a public code repository, and ensuring that all projects are mature and well supported before they join OSGeo are a few of the ways that OSGeo helps ensure longevity of the project and its open code nature.

Another aspect to encouraging the use of OSGeo software is by acting as a marketing department for projects that, for the most part, have no budget for marketing.

Other domains of OSS have some of the same marketing difficulties as geospatial software. Advertising and face-to-face promotion is quite difficult unless you have an organized group, funding and opportunities to meet people. Much is done with an online presence, sharing information with colleagues, and in speaking opportunities at conferences. However, proprietary product marketing goes well beyond that to promote a brand and to put products in the hands of potential customers. Marketing serves to announce the viability of projects across domains and around the world, making it an important part of supporting healthy and more sustainable software development.

Some of the typical marketing activities that OSGeo undertakes include: i) purchasing booth space at tradeshow events; ii) producing material for distribution such as brochures; and iii) organizing workshops or conferences to address regional needs. These all require some sort of financial support and organization. OSGeo's major annual conference event, FOSS4G (http://foss4g.org), moves around the world and is expected to draw up to 1,000 attendees this year. This opportunity gives skeptics and advocates alike the chance to hear about progress and experience hands-on workshops with the tools they are specifically interested in.

Becoming a sponsor is a significant way for a user or group to give back to projects. While OSGeo has project-specific sponsorship programmes, the main funding for OSGeo's activities is through foundation level sponsors that do not target a specific project (http://osgeo.org/sponsorship). This in turn feeds more education and awareness of the projects, building more users and encouraging more contributors.

Previously, it was impossible for organizations and individuals to donate funds for two reasons. First, no legal entity existed to accept the funds. Second, there was no guarantee that the projects supported would survive over the long term. In some cases, sponsors may never have the ability to contribute code or staff time to OSGeo or its projects, but now they can easily provide funds to show their support in a tangible way.

Ecosystem as a Community Aggregator

There is a far-reaching social aspect to the development of OSGeo that also deserves mention. In today's economy, having a community for a product is a gold mine. Product placement has put an endless stream of consumer brands before our eyes: in fast food packaging, in the media, even on children's toys. Developing social networks around these products is ever easier using online tools, encouraging the consumption of products. The cost and planning for pursuing such campaigns must be enormous, though it can be assumed to have enough return on investment to justify continuing the practise.

OSGeo did not pursue this kind of synthetic community development as it happened quite naturally and with little or no cost. In a way, the existence of a productive social network of open source geospatial software users was an impetus to start OSGeo. The number of users of any given piece of software seemed high, but the aggregated number of users across several projects was much higher. Inviting these groups to join together under the OSGeo umbrella provided a new synergy.

Of course, bringing together disparate groups of loosely coupled software users does not guarantee long-term cohesion.

OSGeo users and developers already had a high level of cohesion and this has increased over the past three years. At the most basic level, they were all working from the open source development model. Further below the surface, there were many other overlaps between user groups, developers and even code-level dependencies. Many projects were already working together or building on top of each other.

Just as OSS got a boost by having open source operating systems available, open source geospatial projects encourage each other to develop. These collaborations and dependencies do not begin and end with OSGeo. OSGeo works alongside many other open source projects as well. There are overlaps between several projects in this ecosystem but the playing field is still relatively small, at least compared to the broader open source operating systems. One expected outcome of OSGeo is that by working closely together, there will be more focus on improving existing software rather than starting new projects that overlap others.

By starting an organization where several projects and their members can interact toward common goals, OSGeo tapped into an existing cohesive network of users, developers and organizations. The net effect is that business, research, education, government and more are members of the OSGeo ecosystem. The ecosystem has helped to bring these groups closer together to support one another and to seek new ways of working together.

One simple example that shows the cross-pollination of technologies and businesses that support them is found in the OSGeo Service Providers directory (http://osgeo.org/search_profile). This directory allows potential users to find support in their region, their language, or for the specific type of software they use.

In many cases, the directory shows service providers supporting several different software packages across various geographies and languages, thus supporting the concept of a healthy, diverse, ecosystem. This directory also gives organizations an opportunity to submit their information and make their name known.

Summary

Coupling the development of a social ecosystem with the general goals of providing stability, marketing, and shared resources has helped make OSGeo an open source development success story. It brings together code, users and funding in a way that encourages further growth of solid products in an increasingly competitive business environment.

Removing barriers and finding common goals has helped to move forward not just geospatial technologies, but open source products in general. This provides more options to global organizations, helping them to avoid the risks of proprietary lock-in and black-box business services.

Tyler Mitchell is the Executive Director of OSGeo. He is also the author of Web Mapping Illustrated: Using Open Source GIS Toolkits (http://oreilly.com/catalog/9780596008659/). He has 13 years of GIS experience, much of which involved open source technologies. He can be found speaking at open source and geospatial events around the world and is dedicated to introducing great tools to great people.

"There is one thing stronger than all the armies in the world, and that is an idea whose time has come."

Victor Hugo

The availability of geospatial data sets is exploding. New satellites, aerial platforms, video feeds, global positioning system (GPS) tagged digital photos, and traditional geographic information system (GIS) information are dramatically increasing across the globe. These raw materials need to be dynamically processed, combined and correlated to generate value added information products to answer a wide range of questions.

This article provides an overview of OMAR (OSSIM Mapping Archive, http://www.ossim.org/OSSIM/OMAR.html) web based geospatial processing. OMAR is part of the Open Source Software Image Map (OSSIM, http://ossim.org/) project under the Open Source Geospatial (OSGeo, http://www.osgeo.org/) Foundation. The primary contributors of OSSIM make their livings by providing professional services to US Government agencies and programs. OMAR provides one example that open source software (OSS) solutions are increasingly being deployed in US government agencies. We will also summarize the capabilities of OMAR and its plans for near term development.

Introduction

OMAR is a web based system for archival, retrieval, processing, and distribution of geospatial assets. Satellite and aerial images, vector sets, unmanned aerial vehicle (UAV) video sets, as well as user generated tags and reference items can be easily searched and manipulated with the system. Searching can be performed on the basis of location, time, or any combination of the stored metadata.

OMAR is unique in its ability to dynamically process raw materials and create value added products on the fly. Imagery is orthorectified (geometrically corrected), precision terrain corrected, and histogram stretched on demand. OMAR combines, fuses, or chips areas of interest according to the users needs. Geospatial assets can then be manipulated, viewed, and processed to provide a wide range of value added products. These products are delivered through several mechanisms:

- a planned generated product that is distributed through ftp or email
- results generated with a simple browser interface
- open standards and interfaces such as the Open Geospatial Consortium's web map service (WMS, http://www.opengeospatial.org/standards/wms), web feature service (WFS, http://www.opengeospatial.org/standards/wfs), and web coverage service (WCS, http://www.opengeospatial.org/standards/wcs), as well as tiling services (TileCache, http://tilecache.org)

OMAR is under active development through US Government funding. OSSIM and OMAR development is being funded by a number of intelligence and defense agencies including the Department of Defense, the National Reconnaissance Office, and the National Geospatial Intelligence Agency. While there are plugin modules that are classified, the OSSIM project is managed and maintained on the Internet in an unclassified environment.

OMAR integrates several OSS solutions to provide an online dynamic processing solution.

OpenLayers (http://openlayers.org/) for dynamic mapping, PostGIS for the Postgres database (http://postgis.refractions.net/), the GRAILS (http://www.grails.org) web framework, and OSSIM are a few of the technologies that are used.

The Open Source Business Model

Since 1999, OSSIM has evolved through US Government funding from the Defense and Intelligence communities. Throughout that period, the core development team has worked in a number of different companies while maintaining a close collaborative relationship. The combined OSSIM team has supported a number of different customers. As a result, OSSIM is now deployed in a number of critical government and commercial applications. Over time, a number of solutions and applications have evolved out of the core libraries. Solutions include:

- command line applications
- the ImageLinker prototyping tool (http://www.ossim.org/OSSIM/Downloads.html)
- ossimPlanet for 3D visualization and collaboration (http://www.ossim.org/OSSIM/Downloads.html)
- OMAR for online geospatial management and production

OSSIM has been supported by a number of government agencies through the funding of professional development services. Typically, an agency will hire OSSIM developers to add functionality and meet agency requirements through the use of OSSIM solutions. When viewed in comparison to typical government projects, this approach appears small and fragmented.

However, one of the demonstrated advantages of an open source approach has been frictionless collaboration between all of the projects and contributors. Government customers rapidly become converts to this approach as they begin to inherit and apply improvements that were funded from other agencies and projects. All participants share in the benefits when they contribute with funding support.

The advantages of an OSS approach are slowly becoming evident within the US Government. Changes in policy and acquisition practices are in work to spread the adoption of these practices. Successful projects such as OMAR and OSSIM are providing useful pathfinders as projects experiment and evaluate this approach.

The OSS business model is in the experimental stage within the US Federal government. The agencies are funding some professional services to extend and support OMAR, OSSIM and other OSS geospatial technologies. This model works well for organizations that manage prototype or development projects. Many agencies and organizations within the government do not directly support software development. Front line military operations, known as operational agencies, want certified, configured and supported solutions for their missions. Operational agencies tend to shy away from funding development efforts and a packaged product bundled with maintenance and support turns out to be a better fit for those types of agencies. However, it is becoming clear that open source solutions and support is being implemented across all levels of the government enterprise.

Given the level of interest in and exposure to OMAR, it is clear that a more traditional development team and support staff will be funded to meet operational and developmental requirements.

Government mission requirements have identified a pressing need for additional functionality, an improved user interface, installation packaging, testing, documentation and support.

OMAR Capabilities

OMAR's relational database and geospatial processing capabilities allow rapid generation of value added information products from raw information. Delivery of this processed information can take the form of generated products, web browser views, or on demand web based services.

OMAR's cataloging and provisioning capabilities rapidly locate, process, and distribute value added products across the enterprise. OMAR can automatically detect, ingest, and process when new geospatial assets arrive in any of the monitored repositories. While there are a number of systems that can discover and distribute geospatial assets, OMAR is unique in its ability to process those assets into derivative products and services on demand.

The included OSSIM geospatial processing engine can construct image chains that define the functions, parameters, and conversions that are needed to read, re-project, and process the original geospatial assets into value added derivative products and services. Image chains are parameter driven instructions that describe how to build a value added product. These take the form of spec files that can be stored in the database.

Current installations of OMAR are managing millions of imagery and video files. Even though the current release is still considered in beta, the system is being used to find and rapidly view geospatial assets from multiple repositories.

At this point, the primary user interface is through the browser as development begins to expose specific Web Services Description Language (WSDL, http://www.w3.org/TR/wsdl) services and generated custom products.

A few of the current features of OMAR 2.0 are:

- discovery and online viewing of imagery and video
- national and commercial imagery as well as UAV video is actively ingested, stored and viewed in the system
- online browsing of imagery provided by OpenLayers and OSSIM
- playback of video clips currently provided by an external streaming service
- searching based on location, acquisition time, and metadata
- interfaces with the embedded Post GIS/Postgres relational database provide the ability to select assets based on location, time, or any of the values in the metadata including sensor types, target identifiers, and various collection criteria

WMS Services are currently provided by embedded MapServer (http://mapserver.org/) functionality. The WMS interface coupled with the background OSSIM processing enable products and views to be composed on demand. Initial WFS support is provided through MapServer and OpenLayers. This support will be streamlined and enhanced in future releases.

OSSIM and Geospatial Data Abstraction Library (GDAL, http://www.gdal.org/) provide support for a wide array of native geospatial formats.

These libraries can provide file format conversion and allow assets to be referenced and used without intermediate conversion.

The OSSIM library provides on the fly reprojection and orthorectification (http://en.wikipedia.org/wiki/Orthophoto) for display and output. Satellite and aerial images are processed through OSSIM sensor models to map projected products.

Rigorous and Rational Polynomial Coefficient (RPC) sensor models apply precision terrain correction to Digital Terrain Elevation Data (DTED, http://en.wikipedia.org/wiki/DTED), Shuttle Radar Topography Mission (SRTM, http://en.wikipedia.org/wiki/SRTM), or raster elevation data sets that are referenced to the system.

OSSIM creates complex products from spec files. The spec files define the parameters and processing steps needed to build a product. These image chains are being used in OMAR to define processing and views based on demand. Future development will include the ability to store user defined processes in the database for custom products.

Multi-image mosaics and fusions (http://en.wikipedia.org/wiki/Sensor_fusion) are being produced by the OMAR system in pre-defined image chains. New services and variations will be exposed in the future.

OSSIM and GDAL provide an extensive number of file format conversion services. The user interface needs to be extended to expose this for user generated products.

OSSIM dynamically re-projects data into geographic views. User selectable map projections and datums need to be exposed through the user interface.

The underlying conversion process already exists. User defined areas of interest can be identified through the WMS interfaces. The underlying architecture has the ability to generate area of interest products and services for interfaces with external systems and user requests.

Much of the current development work is focused on exposing existing processing capabilities through the user interface, defining and implementing services, and developing targeted image chains for specific functionality.

OMAR Walkthrough

Access through the system is authenticated through a user login, as seen in Figure 1. LDAP authentication is currently in development. Roles and privileges are granted based on the results. The current version of OMAR supports two levels of access: users and administrators. Once logged in, the user will be presented with an initial interface for browsing and searching.

Typically, the user will search for geospatial assets by selecting a geographic area of interest and filtering by acquisition date, sensor type, target identifiers, or any combination of the metadata tags that are stored in the internal database.

The administrator has the ability to add new tags and criteria to the search panel shown on the left in Figure 2.

The imagery search web interface, shown in Figure 3, provides a map that the user can pan or zoom to select a desired area of interest. Zooming the map to a particular area of interest reveals outlines of data sets that are available in the system. The user can then select the "Area of Interest" mode and draw a selection rectangle over the desired area for search.

Figure 1: OMAR Login screen and Initial Interface

Figure 2: OMAR Web Interface for Searching Geospatial Assets

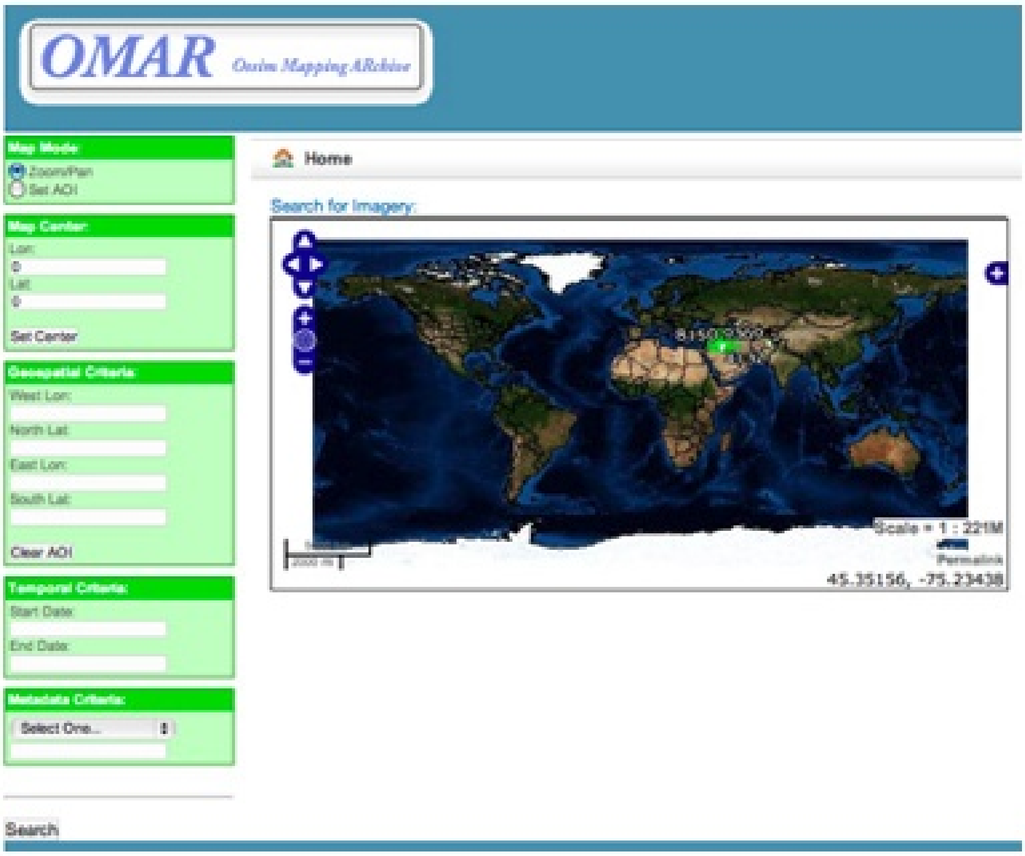

Figure 3: OMAR Imagery Search Interface

Additionally, the fields in the left panel can be filled to further filter the search for data. A number of parameters are available for search criteria. The user can manually enter center and corner coordinates, acquisition time and date parameters, or values for any of the metadata tags. In many government applications this will include sensor ids, target identifiers, sensor types, or resolution criteria. Configurable metadata parameters and overviews of the data assets are displayed. Clicking on the overview thumbnail will allow interactive viewing of the full data set.

OMAR provides interactive zooming and panning into a satellite image. Behind the scenes, OSSIM is processing the raw file through a sensor model, cropping and zooming into the image, and enhancing the image with histogram stretching. Roaming, panning and zooming is accomplished interactively.

The underlying image is being projected through a sensor model, orthorectified, precision terrain corrected and histogram stretched on the fly.

UAV Video

OMAR also can process UAV Predator (http://www.airforce-technology.com/projects/predator/) feeds. These feeds can be searched, selected and played back through the web browser.

OMAR is able to extract and parse the metadata from Motion Imagery Standards Board (MISB, http://www.gwg.nga.mil/misb/) compliant video streams. Missions, acquisition dates and times, platform and center of interest coordinates are used to populate the internal database and position the data for geographic searches.

Summary

OMAR is an integration of several successful OSS projects to provide an enterprise solution for geospatial data management, production, and distribution. Defense and Intelligence agencies of the US government have provided funding to support OMAR and the underlying OSSIM software libraries. Through this support, the OSSIM development team has been employed through a number of collaborating projects.

Interest in OMAR is gaining across a number of agencies and it will soon evolve to a more formal government project. Maintaining OSSIM as an unclassified open source project on the Internet has been key to its success and its ability to collaborate across a number of separate government projects.

The OSSIM team is always looking for additional contributors, developers, and users. Additional information can be found at http://www.ossim.org.

Recommended Resource

OSSIM Whitepaper
http://ossim.telascience.org/ossimdata/Documentation/OSSIM_Whitepaper.pdf

Mark Lucas has pioneered efforts in OSS development in remote sensing, image processing and GIS. Mark established http://remotesensing.org and has led several government funded studies and development efforts since 1996. These efforts include OSSIM projects for the National Reconnaissance Office, the Open Source Prototype Research and Open Source Extraordinary Program projects for the National Geospatial-Intelligence Agency. He leads the Open Technology Development effort within the Department of Defense Advanced Systems and Concepts in collaboration with National Information Infrastructure and the Business Transformation Agency. Mark has a BS in Electrical Engineering and Computer Science from the University of Arizona and a MS in Computer Science from West Coast University. He is on the Board of Directors of the Open Source Geospatial Foundation, the Open Source Software Institute, and the National Center for Open Source Policy and Research. Mark is currently a principal scientist at RadiantBlue Technologies Inc. (http://www.radiantblue.com/).

Scott Bortman is the system architect and primary developer for the OMAR web processing system. He has been a primary contributor to the OSSIM software baseline over the last decade. He has a BS and MS of Computer Engineering from the Florida Institute of Technology and has worked for a number of government contractors including Computer Science Innovations, ImageLinks Inc., Intelligence Data Systems, and L3 Corporation. He is currently working as a Senior Software Engineer for RadiantBlue Technologies Inc. Scott has a strong background in database design, Image Processing, C++ and Java programming. Within the OSSIM development team, Scott is known for his ability to stay current with the latest advances in software development tools, methodologies, and approaches.

"The only way to discover the limits of the possible is to go beyond them into the impossible."

Arthur C. Clarke

Soon after Codd (http://en.wikipedia.org/wiki/Edgar_F._Codd) wrote his paper on relational algebra in 1970, relational databases significantly changed the way people managed data. Today, relational databases are the workhorses of enterprise data storage. Similarly, imagine a world without email or the Internet. What will the next “killer app” or “killer service” look like? What kinds of attributes and features will it provide?

In this article, we provide a primer on geospatial technology. We then explain possible reasons for growth in the geospatial industry, examine Ingres' geospatial project, and relate the material to learnings about open source as a protocol for business.

The Storm is Coming

Technology change has made spatially aware applications and devices more affordable and accessible. This is based on smaller, faster, and more power efficient chips. Increased network bandwidth for both wired and wireless networking has improved the availability of spatial data. Traditionally dominated by a few large competitors, new standards and competition have started the geographic information system (GIS) industry's evolution towards becoming main stream. More importantly, these standards and technologies have hastened the inclusion of spatial awareness into applications from other industries. New opportunities are emerging to add maps and spatial awareness to enterprise information technology (IT) and to do so at a cost that the masses can afford. The stage is set to provide new insights from existing data. As countless more devices become spatially aware and interconnected, we are going to experience an epic storm of spatial data.

Geospatial Data

There are two types of geospatial data: raster data and vector data. Raster data are essentially pictures, although not always in the visible spectrum. Satellite or aerial images are examples of raster data. Vector data are a mathematical representation of real life. Vector constructs include points, lines, polygons, and other shapes which can be used to represent houses, roads, rivers, parks, lakes, and more.

The industry standards were published by the Open Geospatial Consortium (OGC, http://www.opengeospatial.org/) and describe how raster, vector, and combined map data can and should be represented. These standards include Web Coverage Service (WCS, http://www.opengeospatial.org/standards/wcs), Web Feature Service (WFS, http://www.opengeospatial.org/standards/wfs), and Web Map Service (WMS, http://www.opengeospatial.org/standards/wms) which describe serving raster data, vector data, and maps respectively. OGC also defines how relational databases should store and provide interfaces to act upon spatial data. Adhering to these standards means systems can interoperate more easily. This enables using raster data from one source, vector data from another, and combining them into a map service that can be consumed by a large choice of software.

Coordinate Systems

Most of us are familiar with the concept of latitude and longitude with zero degrees longitude centered on Greenwich, England. There are other systems that have zero degrees centered on Moscow, Paris, and other major cities. Each of these systems is a coordinate system.

Since geospatial data may be stored in any one of a number of coordinate systems, it is important to be able to convert between them. The Open Source Geospatial Foundation (OSGeo, http://www.osgeo.org/) sponsored software projects Proj.4 (http://trac.osgeo.org/proj/) and csmap (http://trac.osgeo.org/csmap/) provide this functionality. Another name for a coordinate system is a spatial reference system.

Geodetics

We have all been told that the closest distance between two points is a straight line. But on the surface of a sphere, that straight line is actually an arc. To complicate things further, most planets are not perfect spheres but ellipsoids with imperfections. The science of geodetics (http://en.wikipedia.org/wiki/Geodetics) deals with the measurement of the earth.

Why Use a Relational Database for Spatial Data?

There are a number of formats for storing spatial data, including several that are just files on a disk. So, why burden oneself with the overhead of a relational database? With one user, one set of data, and fairly simple and unchanging demands for data, it is easy to make the case for storing data as files on a disk. However, once you need to share that data with a team of people, things become more complex. A relational database management system (RDBMS, http://en.wikipedia.org/wiki/RDBMS) provides atomicity, consistency, isolation, and durability (ACID, http://en.wikipedia.org/wiki/ACID). In short, this means that the database will ensure that your data is not corrupted. An RDBMS also provides a client/server architecture that allows shared data over a network. The security model of a RDBMS enables roles defining who can view, modify, or delete data. These are all important considerations when sharing data.

Ingres

Ingres (http://www.ingres.com/) was one of the original RDBMSs and was born out of the INGRES project at the University of California, Berkeley in the 1970's. In 1980, INGRES project founders Michael Stonebraker and Eugene Wong created Relational Technology Incorporated (RTI) based on the technology. RTI changed names to Ingres Corporation and was purchased by Ask Corporation in 1990. Computer Associates acquired Ask in 1994. In 2005, Ingres was spun out of Computer Associates with venture funding to form the current Ingres Corporation. Today's Ingres Corporation is an open source startup based in Redwood City, California. Ingres' revenues have recently grown to $68M, despite the gloomy economy, making it currently the largest independent open source RDBMS company.

Ingres competes with closed source offerings from Oracle, IBM, and Microsoft. The main open source RDBMS projects are MySQL (http://www.mysql.com/), now owned by Sun Microsystems, and PostgreSQL (http://www.postgresql.org). O'Mahony and West (http://pascal.case.unibz.it/retrieve/2814/rp-omahony.pdf) propose there are two major types of community, those that are grass roots initiated and those sponsored by a for-profit firm. In the context of the Ingres community today, the latter is a better fit.

Hindsight is 20/20

It is worth noting that Ingres was one of the first RDBMSs to support geometry datatypes. Geometry datatypes provide mathematical constructs to describe points, lines, polygons, and other data types for describing objects and relating them in cartesian space.

Many of these constructs are used to enable geospatial technology to relate objects on the surface of the earth. Even though Ingres supported geometry types, it had no support for coordinate systems, geodetics, and its geospatial functions were sparse. As the industry defined standards for additional data types and functions in the late 1990s, work was needed to update the code to support them. When the Ingres Spatial Objects Library (SOL) was originally developed, the decision was made to outsource its development. The deal left the intellectual property (IP) in the hands of the outsourcing company, leaving Ingres with the rights to distribute binaries, but not the code. Recall that in those days, geospatial technology was a tiny niche and only those with deep pockets and an urgent need for the technology were interested.

Ingres Geospatial Project

Ingres' customer base of over 10,000 customers represents a considerable amount of data and business. Since IT systems often contain spatial data in the form of addresses, it is common for customers and the community to ask what the company is doing in the area of spatial technologies. As an open source company, it is a significant problem to have an in-demand component not available as open source. Out of customer and community interest and the emergence of new standards, the Ingres geospatial project (http://community.ingres.com/wiki/Ingres Geospatial) was born.

Power of Open Source

In IT Doesn't Matter (http://web.njit.edu/~jerry/CIS-677/Articles/Carr-HBR-2003.pdf), Carr notes that large IT suppliers such as Microsoft, Oracle, and IBM are making huge amounts of money while companies overspend on IT.

Carr also notes that there is no correlation between IT spending and superior performance. If anything, the relationship is the inverse. Carr asserts that IT can be done more efficiently and inexpensively as there is no strategic advantage to paying more for platform software. Open source software (OSS), which is distributed for free and has development costs spread across numerous firms, seems well positioned as a commodity and poses a significant threat to the business of the closed source market share leaders.

The success of OSS projects such as Linux, Apache, and Firefox demonstrate that OSS can compete and be successful. In many cases, it can even challenge the market leaders.

To Make a Change, First Look in the Mirror

As a code base that was recently reopened, the Ingres open source community struggled to compete with the enormous mindshare of MySQL. Much like the battle of VHS and Betamax, community developers did not seem to pay much attention to details of how Ingres was technically superior. It is fair to characterize Ingres' early days of returning to its open source roots as "open code" but closed in other ways. While an archive of the source code was available from the website, design discussions, code inspections, the production code repository, product roadmap information, and more were hidden behind the corporate firewall. It does not make sense to be an open source company without benefitting from an open source community.

Changes needed to make Ingres more open to community participation met with resistance within the company. A company exists to "maximize shareholder wealth" by making a lot of money. Open source and making money are not at odds. However, in order to make money with open source, you must first invest. For a startup with a sharp focus on profitability, it is very difficult to set aside money and people to work on something that may not generate a short term return. Despite the odds, a decision to forge ahead was made and investment in infrastructure such as a public code repository, bug tracking system, public technical documentation, community mentorship, and community management were made.

Survey of Reusable Components

It is worth explaining that much of our underlying technology, its defining points, lines, polygons and the functions for operating on them, is a commodity. We call this a "geometry engine" for the sake of this article. Given the importance of community, it was important to look first at existing communities and code reuse. Top on the list of priorities was to contribute to making an existing code base stronger rather than creating yet another geometry engine. Contact was made with members of the OSGeo community who assisted in identifying candidates for code reuse. The leading candidate was a project called Geometry Engine Open Source (GEOS, http://trac.osgeo.org/geos/) originally developed by Refractions Research to enable PostGIS (http://postgis.refractions.net/), the geospatial plugin for PostgreSQL. GEOS had roughly 20 year's worth of investment borne mostly by Refractions. A plan was assembled where Ingres would adopt GEOS and contribute to the development of the code. Helping to make this proposition more attractive, Ingres and others lobbied other companies in the OSGeo community to join. Eventually, the GEOS project was moved to OSGeo, and code contributors came forward from a number of companies. Each of the organizations involved benefits from giving a little to the development of GEOS and receives much more back in return.

OSS provides many benefits including sharing costs, risks, and ideas. OSS enables swift development, open communications, and collaboration. With closed source, just negotiating the legal agreements between the multiple companies involved can take many months. With open source, new companies and people can join the project without having to renegotiate contracts, thus reducing transaction costs.

Summary

The geospatial industry is poised for tremendous growth as location aware applications and devices grow in popularity. Enterprise IT will discover new value and insights through spatial analysis of existing data. Open source can reduce the transaction costs of technology partnerships. Businesses should seek out partners with interests that align through mutual investment and reuse of OSS. Doing so allows them to re-allocate spending to areas that provide unique value.

Andrew Ross is a Director within the Engineering team at Ingres where he leads a team that works on Ingres community projects including Geospatial and CAFÉ. Prior to joining Ingres, Andrew was an Architect and software developer at Nortel. Andrew has been developing and using open source for over a decade and teaching University classes using open source since 2004. He is a charter member of The Open Source Geospatial Foundation. Andrew is Founder and President of the non-profit Free and Open Source Software Learning Centre (http://fosslc.org/).

"Network effects from user contributions are the key to market dominance in the Web 2.0 era."

Tim O'Reilly

In "What is Web 2.0?" (http://www.oreillynet.com/pub/a/oreilly/tim/news/2005/09/30/what-is-web-20.html) Tim O'Reilly identifies the characteristics that distinguish Web 2.0 from the first generation of Web applications. One key aspect is participation. Instead of users simply consuming information, Web 2.0 technology enables all of us to participate in building content. The power of Web 2.0, in Tim O'Reilly's words, is that it provides a platform for "harnessing collective intelligence". Perhaps the best known example of this is Wikipedia, which is distinguished from other online encyclopedias by the fact that its content is provided by users rather than a small group of experts. This model has been so successful, even the Encyclopedia Britannica has adopted a Web 2.0 approach (http://www.infoniac.com/breaking/encyclopedia-britannica.html).

A critical challenge to participation is interoperability--integrating the islands of technology that characterize most information technology (IT) organizations. There have been earlier attempts to create a standard framework for distributed computing such as (CORBA http://en.wikipedia.org/wiki/Corba) and DCOM (http://en.wikipedia.org/wiki/Distributed_Component_Object_Model), but the complexity of these environments has limited their adoption. A more recent and simpler approach is Representational State Transfer (REST, http://en.wikipedia.org/wiki/Representational_State_Transfer). In this article, we begin with an examination of the critical challenges facing organizations responsible for maintaining our utility, telecommunications and transportation infrastructure, outline how open standards are helping to address these challenges, and then discuss how geospatial data and services can be exposed over the Web. We introduce REST, outline a RESTful implementation of geospatial Web services that provides simple and open access to geospatial data over the Web using standard Web protocols, and describe a prototype web site developed using RESTful Web services by the City of Nanaimo.

Critical Challenges

One of the most serious challenges facing organizations responsible for managing infrastructure, including water, waste water, power, gas, telecommunications, roads, and highways, is increasing the productivity of the field force. This challenge has become particularly urgent in North America where, as a recent study (http://www.us.capgemini.com/PlattsStudy/) of the power utility industry documented, industry is facing the problem of an aging field force. Within the next few years, half of the field force, with their deep knowledge of network facilities, will retire to be replaced by young, inexperienced workers. In some sectors, the situation is dire. We recently chatted with an employee of an Arizona utility who said that 50% of the work force at his firm is eligible to retire this year. This aging work force represents a huge loss of collective intelligence. The challenge for these organizations over the next few years is to transfer the knowledge about the network infrastructure currently resident in the heads of experienced, and soon to retire, field workers into the organization's collective knowledge base. Only then can the collective intelligence be harnessed by all workers, and most critically, younger workers, to improve productivity in the future.

Another critical challenge is interoperability. For example, organizations with an engineering focus typically have islands of technology such as CAD (http://en.wikipedia.org/wiki/Computer-aided_design), mobile, GIS (http://en.wikipedia.org/wiki/Geographic_information_system), and tabular financial and business systems. Many of these systems are proprietary, often legacy, developed by different vendors, and are incompatible with each other. Productivity and efficiency are the business forces which are forcing IT organizations to look for ways to break down interoperability barriers.

Open Standards

One of the most important technical advances to provide a foundation for interoperability is open standards (http://en.wikipedia.org/wiki/Open_standards). The Web, which has become the world's operating system, is based on standards from the IEEE (http://ieee.org), IEC (http://www.iec.ch/), W3C (http://www.w3.org), and ECMA (http://www.ecma-international.org/). Geospatial standards from the Open Geospatial Consortium (OGC, http://www.opengeospatial.org/) allow the exchange of spatial data. Web applications from the major geospatial vendors are still for the most part proprietary. But there are open source projects that are moving in the direction of an open Web 2.0 platform.

Shortly after the formation of the Open Source Geospatial Foundation (OSGeo, http://www.osgeo.org/), Autodesk (http://www.autodesk.com) released the source of the Feature Data Object (FDO, http://fdo.osgeo.org/) application programming interface (API) and the MapGuide Open Source (http://mapguide.osgeo.org/) platform to the open source community. FDO is differentent from other programming interfaces. It was designed to support the editing and versioning of spatial data. FDO provides consistent access to a large number of spatial data stores including Oracle Spatial (http://www.oracle.com/technology/products/spatial), SHP (http://en.wikipedia.org/wiki/Shapefile), ArcSDE (http://esri.com/software/arcgis/arcsde), SDF (http://en.wikipedia.org/wiki/SDF) and GDAL/OGR (http://gdal.org/ogr) as well as open standards like KML (http://en.wikipedia.org/wiki/Keyhole_Markup_Language) and WFS (http://www.opengeospatial.org/standards/wfs). The C++ code for FDO, which is available from OSGeo, has been compiled to run on Windows and Linux. Similarly, the source code for MapGuide Open Source is available from OSGeo.

But there is still the issue of how to expose these applications in a general way on the Web. For example, you can wrap a PHP (http://www.php.net/), JSP (http://java.sun.com/products/jsp/), or ASP (http://www.asp.net/) programming interface around an application with a Javascript client, but this approach will be different for each Web application. A more general approach is to wrap the C++ code with standards-based Web services. This not only allows client applications to access geospatial data and services in a standard way, but allows geospatial data and services to be integrated with other Web services using orchestration such as Business Process Execution Language (BPEL, http://en.wikipedia.org/wiki/Business_Process_Execution_Language).

Earlier attempts to create a standard framework for distributed computing included CORBA and DCOM, but the complexity of these environments has limited their adoption. Two more recent approaches are: i) the W3C's Simple Object Access Protocol (SOAP, http://en.wikipedia.org/wiki/SOAP_(protocol)) which is supported by application development tool makers such as IBM, BEA Systems, and Microsoft; and ii) REST which has been used by Amazon, Google, and others to create interfaces to their Web services.

REST

The term REST was introduced by Roy Fielding in his Ph.D. dissertation (http://www.ics.uci.edu/~fielding/pubs/dissertation/rest_arch_style.htm) and describes an architecture style of networked systems. The motivation for REST has been to rely on the simplicity of the HTTP protocol and data exchange based on XML (http://en.wikipedia.org/wiki/Xml) and MIME (http://en.wikipedia.org/wiki/MIME). Since REST uses standard HTTP methods, RESTful applications are not hindered by firewalls. By linking to open source components such as MapGuide, FDO, and other open source libraries, a geospatial Web services framework, tentatively named king.rest, has been developed (http://www.jasonbirch.com/nodes/tag/sl-king/). It enables a site administrator with no programming experience to deploy HTML, KML, and other representations of geospatial data together with metadata pages that expose information about the Web services provided by the site in a form that is easily crawlable by search engines, and easily understood by anyone wanting to access the data for other applications.

Using the king.rest framework, the City of Nanaimo has implemented a prototype geospatial web services site where all of the City's public data will be exposed through a single URL (http://maps.nanaimo.ca/data/). The first incarnation provides read-only access to the city's geospatial data. However, with an FDO data provider that supports geospatial RESTful Web services, any application that supports FDO will have edit capabilities over the Web.

The index is a static HTML page that allows users to search for data. A street index enables Web crawlers to access every property in the city. Geospatial data in different representations, such as KML, JSON, XML, or PNG, can be accessed simply through a URL.

Summary

RESTful geospatial Web services can provide simple and open access to geospatial data over the Web using FDO and standard web protocols. Because FDO is differentiated from other programming interfaces as being designed to specifically support the editing of spatial data, a RESTful implementation of FDO enables full edit access to geospatial data and provides a Web 2.0 platform that can help address the challenges of the aging work force and interoperability.

Haris Kurtagic is General Manager and Founder of SL-King which focuses on developing spatial applications. Haris received an M.Sc degree in Information Science from the University of Sarajevo. His interests include spatially-enabled relational database management systems and the opportunities offered by open source geospatial. His first open source project was an Oracle Spatial data provider, King.Oracle, for the FDO interface.

Geoff Zeiss joined Autodesk in 1999 where he is Director of Technology. His interest is in geospatially-enabled solutions and infrastructure management. In 2004, Geoff received one of ten annual global technology awards from Oracle Corporation for technical innovation and leadership in the use of Oracle. Prior to joining Autodesk, Geoff was Director of Software Engineering at VISION Solutions which was widely recognized for developing innovative infrastructure solutions using leading edge technology.*

An Empirical Investigation into the Adoption of Open Source Software in Hospitals

Copyright: Gilberto Munoz-Cornejo, Carolyn B. Seaman, A. Gunes Koru

From the Abstract:

Open source software (OSS) has gained considerable attention recently in health care. Yet, how and why OSS is being adopted within hospitals in particular remains a poorly understood issue. This research attempts to further this understanding. A mixed-method research approach was used to explore the extent of OSS adoption in hospitals as well as the factors facilitating and inhibiting adoption.

The findings suggest a very limited adoption of OSS in hospitals. Hospitals tend to adopt general-purpose instead of domain-specific OSS. We found that software vendors are the critical factor facilitating the adoption of OSS in hospitals. Conversely, lack of in-house development, as well as a perceived lack of security, quality, and accountability of OSS products were factors inhibiting adoption. An empirical model is presented to illustrate the factors facilitating and inhibiting the adoption of OSS in hospitals.

http://www.research.umbc.edu/~cseaman/papers/IJHISI08.pdf

April 8

IT360

Toronto, ON

IT360 is where IT professionals and business decision makers get knowledge through in-depth conference sessions. IT360 focuses on unified communications, security, cloud computing and virtualization.

http://www.it360.ca/

April 17-18

COSSFEST

Calgary, AB

COSSFEST demonstrates the use of open source to increase the capabilities and competitiveness of your business. COSSFEST brings together thought leaders in the technology industry, technology professionals, businessmen, and computer enthusiasts to share the latest news and open source know-how.

http://www.cossfest.ca/

May 4-6

MCeTech

Ottawa, ON

MCETECH aims at bringing together researchers, decision makers, and practitioners interested in exploring the many facets of Internet applications and technologies.

http://www.mcetech.org/

May 10-13

CNIE International Conference

Ottawa, ON

With an expected attendance of over 400 national and international delegates working in the fields of educational technology, health education, K-12 education, multi-media design and distance learning, CNIE offers a unique opportunity for learning, networking and idea exchange. Join colleagues from across the education spectrum discussing, debating and exploring the integration of learning and technology.

http://www.learningconference.ca/cnie2009

May 12-13

BSDCan

Ottawa, ON

BSDCan has established itself as the technical conference for people working on and with BSD based operating systems and related projects. The organizers have found a formula that appeals to a wide range of people from extreme novices to advanced developers.

http://www.bsdcan.org

May 13-15

SummerCamp

Ottawa, ON

This event will bring together industry, academia, government, and community to learn about open source and to encourage cross pollination of ideas and talent.

http://www.fosslc.org/drupal/summercamp2009

May 16-17

MSR Mining Challenge

Vancouver, BC

The MSR Mining Challenge brings together researchers and practitioners who are interested in applying, comparing, and challenging their mining tools and approaches on software repositories for open source projects. Unlike previous years that have examined a single project or multiple projects in isolation, this year the MSR challenge involves examining the GNOME Desktop Suite of projects. The emphasis this year is on how the projects are related and how they interact.

http://msr.uwaterloo.ca/msr2009/challenge/index.html

May 16-24

ICSE

Vancouver, BC

ICSE, the International Conference on Software Engineering,® is the premier software engineering conference, providing a forum for researchers, practitioners and educators to present and discuss the most recent innovations, trends, experiences and concerns in the field of software engineering.

http://www.cs.uoregon.edu/events/icse2009/home/

May 17-22

Open Source Programs for Mac

BC Public School System

In this knowWEEK we will look at open source programs for Mac that can be used for browsing, video podcasting, instant messaging, emails, podcasting, video playback, word processing/office suites, sound recording, publishing, as well as others. We will also look at ways that you can use these tools in your classroom and share some examples of how teachers are currently using these in classrooms.

http://knowschools.ca/moodle/mod/book/view.php?id=12

The goal of the Open Source Business Resource is to provide quality and insightful content regarding the issues relevant to the development and commercialization of open source assets. We believe the best way to achieve this goal is through the contributions and feedback from experts within the business and open source communities.

OSBR readers are looking for practical ideas they can apply within their own organizations. They also appreciate a thorough exploration of the issues and emerging trends surrounding the business of open source. If you are considering contributing an article, start by asking yourself:

1. Does my research or experience provide any new insights or perspectives?

2. Do I often find myself having to explain this topic when I meet people as they are unaware of its relevance?

3. Do I believe that I could have saved myself time, money, and frustration if someone had explained to me the issues surrounding this topic?

4. Am I constantly correcting misconceptions regarding this topic?

5. Am I considered to be an expert in this field? For example, do I present my research or experience at conferences?

If your answer is "yes" to any of these questions, your topic is probably of interest to OSBR readers.

When writing your article, keep the following points in mind:

1. Thoroughly examine the topic; don't leave the reader wishing for more.

2. Know your central theme and stick to it.

3. Demonstrate your depth of understanding for the topic, and that you have considered its benefits, possible outcomes, and applicability.

4. Write in third-person formal style.

These guidelines should assist in the process of translating your expertise into a focused article which adds to the knowledgable resources available through the OSBR.

Upcoming Editorial Themes

April 2009:	Open APIs Guest Editor: Michael Weiss, Carleton University
May 2009:	Open Source in Government Guest Editor: James Bowen, University of Ottawa
June 2009:	Women in Open Source Guest Editor: Rikki Kite LinuxPro Magazine

Formatting Guidelines:

All contributions are to be submitted in .txt or .rtf format.

Indicate if your submission has been previously published elsewhere.

Do not send articles shorter than 1500 words or longer than 3000 words.

Begin with a thought-provoking quotation that matches the spirit of the article. Research the source of your quotation in order to provide proper attribution.

Include a 2-3 paragraph abstract that provides the key messages you will be presenting in the article.

Any quotations or references within the article text need attribution. The URL to an online reference is preferred; where no online reference exists, include the name of the person and the full title of the article or book containing the referenced text. If the reference is from a personal communication, ensure that you have permission to use the quote and include a comment to that effect.

Provide a 2-3 paragraph conclusion that summarizes the article's main points and leaves the reader with the most important messages.

If this is your first article, include a 75-150 word biography.

If there are any additional texts that would be of interest to readers, include their full title and location URL.

Include 5 keywords for the article's metadata to assist search engines in finding your article.

Copyright:

You retain copyright to your work and grant the Talent First Network permission to publish your submission under a Creative Commons license. The Talent First Network owns the copyright to the collection of works comprising each edition of the OSBR. All content on the OSBR and Talent First Network websites is under the Creative Commons attribution (http://creativecommons.org/licenses/by/3.0/) license which allows for commercial and non-commercial redistribution as well as modifications of the work as long as the copyright holder is attributed.

The Talent First Network program is funded in part by the Government of Ontario.

The Technology Innovation Management (TIM) program is a master's program for experienced engineers. It is offered by Carleton University's Department of Systems and Computer Engineering. The TIM program offers both a thesis based degree (M.A.Sc.) and a project based degree (M.Eng.). The M.Eng is offered real-time worldwide. To apply, please go to: http://www.carleton.ca/tim/sub/apply.html.