

Bottom Line Pricing in the Land of Oz

Spatial Web Site Architecture

Jonathan W. Lowe

Five years ago, interactive spatial Web sites were rare. A few daring municipalities launched their assessors' parcel map sites with great fanfare while the majority wistfully watched and waited. A few hours surfing today's Web, however, suggests the tide has turned — the current crop of spatial Web sites is both plentiful and sophisticated. Over time, the question has shifted from "Should we do it?" to "How do we do it and at what price?" This column attempts to answer that second question by explaining the basic requirements for putting interactive spatial Web pages online, and comparing the costs and services of several Web hosting companies that cater to spatial customers. The (wide) range of annual costs, depending on the architecture you choose, is between \$300 and \$30,000. We're not in Kansas anymore — choose wisely!

First ruby slipped steps

An interactive spatial Web site is an online map that users can pan, zoom, and query with any Web browser. If your organization decides to begin building an interactive spatial Web site today, the first step will be assembling the necessary puzzle pieces. At bare minimum, interactive spatial



Net Results columnist **Jonathan W. Lowe** covers the role of emerging technologies in the exchange of spatial information. Lowe

is the owner of Local Knowledge Consulting (Berkeley, California), where he designs and implements spatial Web sites. He can be contacted at info@giswebsite.com.

Who's the host with the most? Setting up a spatial Web site requires assembling pieces of the puzzle and an investment of time, money — or both.

Web sites need an IP address, network connectivity, a CPU, data storage space, Web server software, map server software, and a customization language (plus your spatial data, of course). Setting up each of these components of the spatial Web site infrastructure requires an investment of either time, money, or both, and may not be interchangeable. We'll look at each in turn.

Hanging your shingle. For an adoring public to find and view your Web site, it must have an IP address and be connected to the Internet. These two elements, though almost too obvious to merit mention, can pose the largest expense of the overall infrastructure. Unless, that is, you are willing to cooperate with other people.

Because of their graphics-heavy nature, the typical spatial Web site requires a fast Internet connection. A T-1 connection costs more than \$1,000 per month, not including the expense of running the lines to your facility. Even if this is not cost-prohibitive for your organization, consider that even faster connectivity, such as a 42–100 megabit connection, is much less

expensive if you're willing to share. In any metropolitan area, there is at least one company paying the high monthly charges for both a collection of IP addresses and "thick pipes" to the Internet. These companies survive, even prosper, when multiple customers collocate their computers at this one point of connectivity, each borrowing one or more IP addresses, each benefiting from the fast connection, and each paying about \$150 per month for the privilege. On the down side, if any one Web site in the group suddenly gets swamped with hits, everyone else's Web sites will also become sluggish. But if the hosting

facility's managers monitor traffic carefully, this (rare) problem is avoidable or can be mitigated shortly after it begins.

Stocking your shelves.

Colocation requires hardware with one or more CPUs and enough disk storage space for operating system software and spatial data. Though fast computers with preinstalled operating systems and vast reserves of disk space are gloriously inexpensive these days (just \$500 buys a powerful box), maintaining that system for online service is not so cheap. Maintenance

Glossary

ASP: Application service provider

CPU: Central processing unit

DNS: Domain name server

DSL: Digital subscriber line

IIS: Internet information server

IP: Internet protocol

ISP: Internet service provider

OS: Operating system

SQL: Structured query language

XML: Extensible markup language

includes ongoing tasks such as conducting security checks, monitoring uptime, and performing upgrades — a time-sink draining \$200 or more per month.

Still sound too expensive? Don't worry, no need to get off that sharing train yet. Extending the colocation business model to those unable to buy hardware or uninterested in administering systems, many Web hosting companies also provide shared access to their own collocated computers. In this arrangement, multiple users each pay a low monthly fee in exchange for a limited portion of disk space on the same collocated machine. The monthly fee also pays an administrator to keep the system secure and performant. These so-called "shared access plans" usually include Web-based email, FTP access, a secure shell, and a public Web site, and are the standard for many small businesses desiring a simple Web presence.

The monthly charges for this arrangement vary (considerably), usually based on the amount of disk storage space and the data-exchange volumes. For example, a site with 100 static Web pages might fill 25 MB of disk. Each time someone visits a page, the site's computer sends that page's text and images across the Internet to the requesting browser. If the average page contains 100 KB of text and image data, then 10 page hits will meter one megabyte of data-exchange. Over the course of a month at a large and/or popular site, data exchange can swell to gigabyte levels. Shared access plans slice and dice both disk access and data-exchange volumes with the aim of balancing their users-per-computer ratios (and turning a profit).

Who's behind the counter? Because so many businesses now have their own Web sites, plenty of companies offer shared access plans, templates for building new Web sites, and easy migration paths to handle increased demand. Spatial Web sites, however, do not fit neatly into the mainstream Web site business that feeds large providers. That's because interactive spatial Web sites require special map serving software; without considerable demand for the capability, large hosts don't want to support the spatial

needs. So, with the introduction of a spatial component, the list of available Web hosts shrinks dramatically.

Before covering some representatives of the spatial Web hosting business, it's worth understanding the software behind interactive spatial Web pages. All Web sites rely on Web server software, the most popular of which is the free, open-source product, Apache (www.apache.org), with a 53 percent market share. Microsoft's (www.microsoft.com) IIS (at \$270) holds second place with approximately 30 percent of the market. As your site's traffic cop, a Web server receives requests for Web pages, locates the requested data on disk, and returns a copy of that data to the requesting browser.

In concert, Web servers and browsers know how to exchange and interpret text, images, audio, video, and a growing list of other formats such as XML. But reading and drawing spatial data formats, such as shapefiles, is not currently supported. Instead, Internet map server software makers provide two strategies — home-cooking or browser-beef-up.

The home-cookers enhance the Web server with map server software that draws images on demand and embeds them into Web pages. Each browser request for a new zoomed or panned map view is actually just a bounding box and click location that the map server software interprets, converts to a snapshot, and returns as a GIF, JPEG, or PNG file.

The browser-beef-up strategy is the same as home-cooking, but adds plug-in functionality to the browser. The plug-in interprets streams of coordinates, transforming the browser into a desktop GIS tool with increased functionality over pure image-based solutions. For instance, browsers enhanced with spatial plug-ins let users highlight groups of features or pan without refreshing the page.

Autodesk's MapGuide (www.autodesk.com), ESRI's ArcIMS (www.esri.com), Intergraph's GeoMedia Web Map (www.intergraph.com/gis), MapInfo's miAware (www.mapinfo.com), and University of Minnesota's (UMN's) MapServer (mapserver.gis.umn.edu) all generally fit one or both of the above models, serving maps in partnership with

the Web server. All but UMN's free MapServer cost between \$1,500 and \$4,500 per CPU. There is also the cost in time of installation: the process of uniting the two software packages (Web server and map server, for any of these five offerings) is notoriously challenging, requiring a breadth of knowledge about networks, Web servers, and programming languages that is often new territory for the average spatial professional.

Stamping your trademark. Wait a minute! What's this about programming languages? Map serving software usually comes with a template or wizards that create a rudimentary interactive spatial Web site with any collection of spatial data. Once the software is installed, users can quickly establish a page with pan, zoom, and identify tools, all without programming a single line of code. Inevitably, though, a template is inadequate. Typical projects require at least three custom functions. For instance, the specification may demand address geocoding or point-in-polygon functionality. Or the contents of the map may link to the contents of the surrounding Web page's text.

Each different map server product has a different customization strategy, all involving some sort of programming or scripting language(s) — from Javascript to Perl, to Visual Basic, all free except for Visual Basic (\$450–\$699 per license). In the flow of a request, the Web server calls a program (written in, say, Java) that tells the map server to generate an image, which the Java program then rolls into a Web page it creates on the fly. Other programs such as Macromedia's (www.macromedia.com) Cold Fusion or Microsoft's Active Server Pages might also be part of the software mixture (and expense accounting).

Scarecrow, Tinman, Lion

If your desire to create an interactive spatial Web site now teeters in a precarious balance with your limited budget or fear of complexity, don't despair quite yet. Though you might begin alone on the yellow brick road to a spatial Web site, an eclectic collection of travelers will soon burst from the wilderness to join you. Following that road as an experi-

mental user myself, I found four friendly Web hosts that cater specifically to spatial customers. Each has a different philosophy, business model, and pricing structure, but all are small, responsive, and eager to succeed.

If I only had a brain. Some organizations have the data and the application requirements, but lack requisite in-house skills (or possibly budget) to produce a spatial Web site. Latitude Geographics Group, Ltd. (www.latitudegeo.com), a company of 12 full-time spatial gurus for hire in Victoria, British Columbia, Canada, offers a service called Geocortex Internet Mapping for just this need. Their depth of experience in spatial technology was immediately clear during an interview with Steven Myhill-Jones, Latitude Geographics' president and CEO. As unassuming as L. Frank Baum's brainy scarecrow, Myhill-Jones seems in eternal pursuit of the next interesting Internet-enabled capability. (Geocortex's latest advance, for instance, is browser-based 3D terrain modeling.) Though "organically driven by customer need," said Myhill-Jones, Geocortex's current customers all use ESRI's ArcIMS software.

Latitude Geographics has grown as an ASP for the past three years by expanding its offering to pivot on services and transferring knowledge to clients. "People are uncomfortable with hosting on an instinctual level," said Myhill-Jones. "Thanks to debacles like WorldCOM, there are concerns with keeping valuable data elsewhere, and trusting entirely in specialists rather than being self-reliant." Furthermore, many of Geocortex's customers in state and local government have policies forbidding extended use of ASPs.

As a result, the Geocortex service has no long-term contract requirements, often hosting only initially, during the development and knowledge transfer phases of spatial Web site projects. In addition to full hosting, customers can also deliver a computer and data to Latitude Geographics for software installation and application development, then collaborate online to develop the site. When the work meets with satisfaction, the customer can migrate their machine in-house, or continue the cohabitation

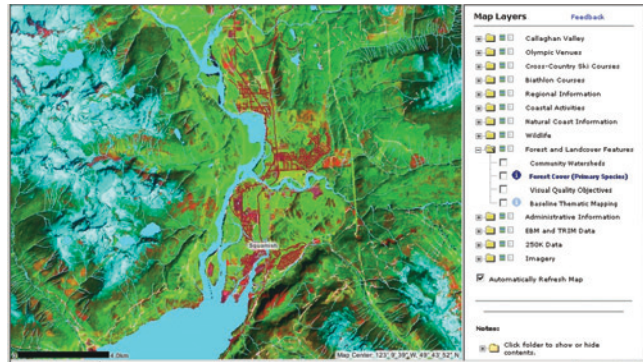


FIGURE 1 The Whistler Corridor Land and Resource Information Server is an example of an online GIS application powered by Geocortex.

arrangement for a monthly fee. Those comfortable with remote hosting can use Web browsers both to alter the application interface, and even to edit their data online, using Geocortex's Internet Mapping Framework software.

Pricing is determined on an ad hoc basis depending on each customer's unique requirements. Generally, projects range between \$300 and \$800 per month, with the most popular package costing \$495 per month, including 24/7 support. To date, Latitude Geographics has not reported any security breaches. A list of Internet reference sites using the Geocortex service appears at www.geocortex.net/gallery (see Figure 1).

If I only had a heart. The sentimental tinman of spatial Web hosting, LMI.net (www.lmi.net), a company of 12 full-times people in Berkeley, California, brings a kinder, gentler approach to technology. In addition to the usual ISP technical offerings, LMI.net focuses heavily



FIGURE 2 Though spatial websites are not their core business, LMI.net's hosting architecture is extensible with UMN's Mapserver and Perl Mapsript, as demonstrated by this City of Berkeley crime mapping application at www.giswebsite.com/cgi-bin/b19.cgi.

on customer satisfaction and affordability (including discounts to local nonprofit organizations). LMI.net is well-suited to Berkeley's culture, openly encouraging local patrons to "Eschew corporate globalization!" by hosting with a small local provider.

Also like most ISPs, the majority of

LMI.net's customers are small- and medium-sized businesses with nonspatial Web sites. Unlike most ISP's however, LMI.net offers users secure shell access to their accounts. This means users can write, compile, and run programs, modify their Web sites directly, and, in short, do anything (within the bounds of security) that the UNIX command line allows.

For a \$50 one-time setup charge and \$25 per month hosting fee, customers get 25 MB of storage space and 2 GB per month of data exchange over their high-speed connection. Spatially, LMI.net's shell access allowed me to install (for free) UMN's Mapserver in conjunction with a customization module called Perl Mapsript. Not a full-fledged GIS, but adequate and speedy enough for most interactive mapping needs, Mapserver proved itself well at mapping the crime density in (where else?) the City of Berkeley, all with self-digitized or free data (see Figure 2).

Though LMI.net does not intend to be a spatial services provider, they are prepared to assist other developers with the installation of their own Mapserver/Mapsript-enabled account. They prove the point that any Web host can become a spatial Web host if willing to install map server software.

If I only had the nerve! Fearful about hanging your valuable data out



FIGURE 3 Metropolis New Media installs their customers' systems in this Northern California data center.

to dry? Courage! The anything-but-cowardly lion of spatial Web hosting, Metropolis New Media (www.metropolisnewmedia.com), is ready to help. A company of two full-time and three part-time people based in San Jose, California, Metropolis' number one concern is security. In fact, secure hosting is Metropolis' only concern, as they refer all spatial services requests to their business partners. For \$1,000 per month, Metropolis scrubs a system with various OS-level cleansing utilities, installs all required Web and map server software, including ESRI's ArcIMS, and colocates the system in a high-speed facility. (see Figure 3). After that, building the Web site is up to you (or your hired help).

To protect their clients' sites, Metropolis analyzes network traffic for any of the thousands of known attacks (such as spoofed IP headers), drops suspicious traffic in real-time if necessary, and alerts appropriate personnel. Unlike a firewall security system where a port is either open or closed, Metropolis uses "stealth-response" algorithms that analyze and control content in both directions, in and out of the server. For instance, if specified, LizardTech (www.lizardtech.com) MrSID file formats can be loaded into the system, but never output.

The company's cofounder and CEO, David L. Hunsinger, explained that (unlike state and local government clients), entities such as banks may have security policies forbidding the use of

Internet services on their network. Consequently, to provide site-selection services to their branches without violating corporate policy, Metroedge, a Chicago-based bank, uses Metropolis as its spatial Web host. Another of Metropolis' well known customers is David Rumsey, whose public map collection is online at www.davidrumsey.com.

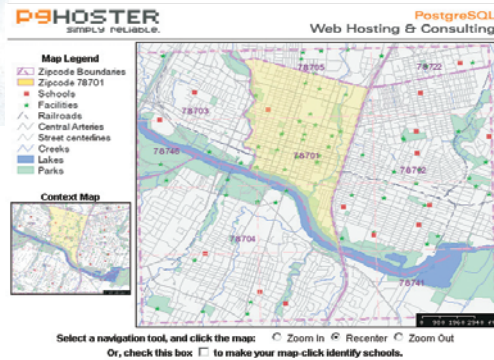


FIGURE 4 pgHoster's online mapping application, demonstrated at www.pgHoster.com/gis, introduces an affordable, analytic, online spatial database.

And Toto, too. At first glance, it would be easy to mistake the shockingly inexpensive, enthusiastic Toto of spatial Web hosting for a too-good-to-be-true prank. But to dismiss pgHoster (www.pghoster.com), a consulting and Web hosting company employing three full-time people in Austin, Texas, would be like ending the Wizard of Oz as Elvira Gulch bicycles away with Toto trapped in her basket. No way, the story's just beginning! In fact, pgHoster may be the spatial community's very first glimpse of an affordable, analytic, online spatial database.

For a one-time setup fee of \$45 plus a \$10 per month hosting charge, pgHoster's spatial customers get an unlimited number of PostgreSQL databases, 300 MB of storage, and 8 GB of data-exchange per month. Included in the setup fee is installation of PostGIS, a free, open-source spatial extension to the PostgreSQL database from Refractions Research (www.refractions.net), as well as UMN's Mapserver and both PHP and Perl Mapsript, for customization. New users will find working templates with sample data ready for substitution with their own data (see Figure 4), which can be loaded from shapefile format directly

into the database. As detailed in the *Geospatial Solutions*, June 2002, Net Results, "Spatial on a Shoestring," the PostGIS-enabled database supports many analytic functions for use within SQL statements. Supported functions include aggregation, querying data for intersection, nearest neighbors, and distance between multiple features.

pgHoster derives its name from the free, open-source database, PostgreSQL (www.postgresql.org). One year after opening its doors to customers requiring online access to an industrial-strength database, pgHoster has acquired close to 1,000 accounts that generate a support load of 35 requests per day on average. Thanks to the fact that its operating system (Linux), database, Web server, and map server software are free, pgHoster can focus on system administration and security, while passing along 100 Mb Internet connectivity at substantial savings to their customers. Though not exclusively oriented toward spatial Web hosting, pgHoster gives traditional GIS vendors a tough act to follow in our increasingly IT-centric spatial industry.

No place like home?

Who knows why Dorothy Gale went home to Kansas instead of trying to bring Auntie Em and Uncle Henry to Oz — after all, that Kansas farm was poverty stricken and the Emerald City was booming! Psychologically, hosting your interactive spatial Web site "from home" may feel right. It may even be company policy. But the monetary bottom line — ultimate ruler of most companies — increasingly favors Web hosting solutions. As this column went to press in April, Intergraph announced its acquisition of Terra Map Server (www.terramapserver.com), effectively extending their offering to include spatial Web hosting. Autodesk referred to California CAD Solutions (www.calcad.com) as an emerging Web hosting provider with MapGuide. ESRI and MapInfo declined public comment. Clearly, nobody really knows just yet how Web services and spatial Web hosting will change our industry, but interesting options continue to emerge. ☺