



OPEN SOURCES #18: THE US ELECTION 2008

KARSTEN VENNEMANN LOOKS AT MAPPING VOTERS FOR THE OBAMA CAMPAIGN WITH MAPSERVER, POSTGIS AND OPENLAYERS

Some months ago I was lucky enough to secure a couple of small contracts with the Obama for America Campaign to support their general effort and the work of their field organizers by providing interactive web mapping capabilities. After some phone discussions with their technology team, it actually was not hard to convince them to implement an extensible and interoperable Web GIS solution entirely based on Open Source software packages. This column outlines the setup of the Obama campaign's Web GIS solution and summarizes the roles of each of the components used.

Goals of the Web GIS

The goal of the Web GIS was to support the work of the campaign field organizers by providing access to up-to-date maps of election trends and information about individual voters targeted by the campaign. Examples of data mapped are: registration percentages of voters, voter persuasion rates, and voting activities of prior elections for each voting district (precinct).

Information about individual voters and their likely candidate of choice were represented as point layers and stored in spatial databases. The initial server for the application would support web access for a few hundred organizers around the country, scattered among up to twenty "battleground" states. The temporary character of the election campaign made it easy to choose the Amazon EC2 computing cloud to host the application because it made the application flexible and scalable for a reasonable price. Another advantage of using the cloud is that it enables fast turnaround times if the application has to be scaled and additional servers are to be needed. Once the complete

software stack is installed on one server, an image can be taken from it and can be made available as the basis for booting up identical virtual machines within a matter of hours, and often within minutes. Using an open source software stack for this purpose gives complete freedom to scale the application to as many servers as needed without running into licensing issues.

Components of an interoperable Web GIS

To start off the initial server we modified an existing Linux-based server image that included most of the Open Source Web GIS application stack already. The EC2 server was running on an Ubuntu (Debian Linux derived) operating system with an Apache 2 HTTP Server front end, PHP for server-side scripting, PostGIS as the spatial database, OpenLayers as a map viewer, and UMN MapServer in CGI mode as the map rendering engine. (*Links to all these projects are provided below.*)

The role of each of the application components is illustrated in Figure 1, "Schematic View of the Interoperable Web GIS". The three main building blocks of the Web GIS are a spatial database (PostGIS), a map rendering engine (MapServer), and a mapping framework which supplies the map viewer (OpenLayers). PostGIS [3] is a spatial extension of the PostgreSQL database [4] and enables GIS functionality such as analysis and intersections via simple SQL queries. MapServer [5] is a CGI application that resides on the server and serves as a map (image) renderer. MapServer can read many different data formats and services such as shapefiles and PostGIS databases. The output results can be easily published as Web Map Services (WMS) via the MapServer configuration (.map) file. OpenLayers [6] is a client side mapping

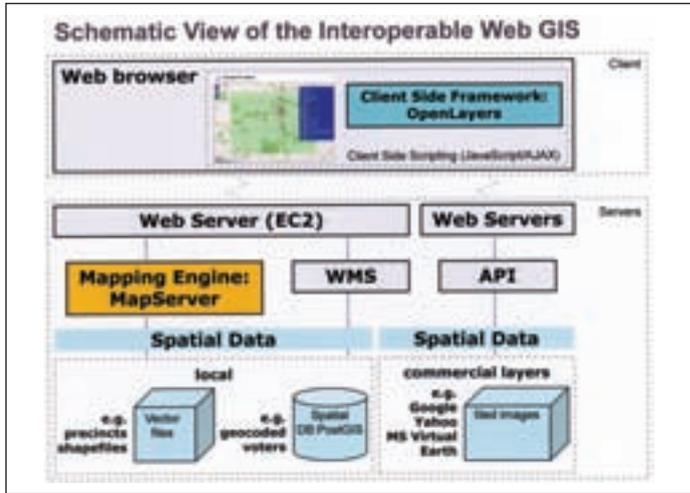


Figure 1. Schematic view of the interoperable Web GIS.

framework/API based on JavaScript/Ajax. The OpenLayers viewer can also take in many different formats and data sources. In our application, OpenLayers enabled the use of commercial map tiles (Google, Yahoo, or MS Virtual Earth) as background maps and provided the functionality to drape the campaign data over it using Web Map Services published by MapServer.

An application template for battleground states

The initial application was set up for the state of Indiana as a template that could be used as a model to build similar applications for other battleground states to follow. The idea was that all battleground states would have their own dedicated website and data sources, but share the server applications such as libraries, the MapServer rendering engine, the PostGIS database and the OpenLayers viewer framework. Two million geocoded voter records were uploaded into PostGIS and could be accessed in the OpenLayers based viewer as a MapServer WMS point layer on top of a selection of commercial background layers, such as Google satellite images, street or terrain layers, or Yahoo street layers. A variety of base layers, such as county and precinct boundaries, along with daily changing information on voter related ratings (as a transparent polygon layer on precinct level) were made available and could be draped over the background layers.

Performance of this "all on one server setup" was reasonable, given that of approximately two million color-coded voter points in the PostGIS database only a portion would typically be displayed at high zoom levels. This approach worked well for the campaign and later was rolled out to about a dozen battleground states.

Another piece of work involved setting up a tool for assigning a list of voters to an specific organizer to contact for campaigning purposes. The setup included a tool for digitizing polygons in the OpenLayers viewer (which is standard functionality) and the ability to customize the tool in such a way that it allows the selection of all voters located within the digitized area via a simple spatial SQL query using PostGIS (via a point-in-polygon query). Subsequently an update of a column in the database with the name of an organizer is triggered by the tool and a second tool enables the export of the respective voter lists from the database via PHP to text files for download.

The Results

Figure 2, "Colorado for Obama", shows the OpenLayers Viewer displaying a Google map street layer in the background with a transparent voter registration precinct layer draped over it; the data is served via WMS and rendered by MapServer from shapefile data. Figure 3, "New Mexico for Obama", again shows a Google map street layer but with voter points draped over it; WMS is used, and the image is rendered by MapServer from PostGIS data. (Note the actual maps used in the campaign are equivalent to what is shown here; the exact data has been altered to protect confidentiality.)

From a technology standpoint, the whole project was a simple use of interoperable and readily available open source geospatial software

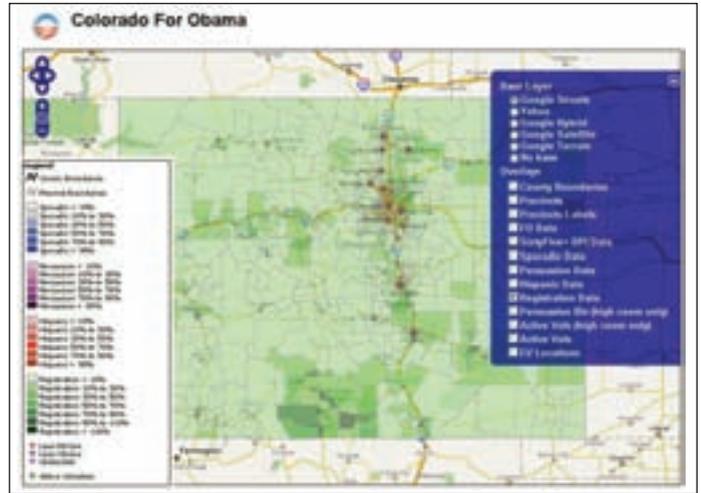


Figure 2. OpenLayers Viewer of "Colorado for Obama"

components. This use case illustrates how powerful and interoperable those components can be, even with minimal customization and on a low budget. For further reading on the current state of open source GIS and web-based mapping, I would suggest Emanuel Schütze's thesis [7] and Paul Ramsey's whitepaper [8].

Acknowledgements:

Thanks to Josh Livni of Umbrella Consulting who created the public image and who also has instructions on his website on how to set this up quickly [1] and to Eric Hammond who created the original EC2 image which can be found at [2].

Links:

- [1] os.umbrellaconsulting.com/wiki/AmazonEC2
- [2] developer.amazonwebservices.com/connect/entry.jspa?externalID=1065
- [3] postgis.refractory.net
- [4] www.postgresql.org
- [5] ms.gis.umn.edu
- [6] www.openlayers.org
- [7] "Current state of technology and potential of Smart Map Browsing (2007) in web browsers using the example of the Free web mapping application OpenLayers" (Emanuel Schütze). www.smartmapbrowsing.org/thesis_EmanuelSchuetze.pdf
- [8] "The State of Open Source GIS, Version September 2007" (Paul Ramsey). www.foss4g2007.org/presentations/viewattachment.php?attachment_id=8

Author's Bio

Karsten Vennemann (karsten@terragis.net) is a GIS professional with a background in Geography and Soil Science. His work mostly focuses on GIS in the context of natural resources, sustainable development and social justice. In 2007 he started his own company, Terra GIS (<http://terragis.net>), based in Seattle, Washington, USA.

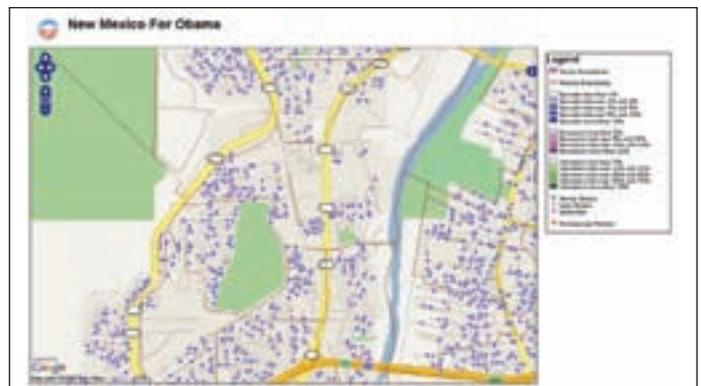


Figure 3. Google map street layer of "New Mexico for Obama"