

# Easy Web Based GIS Analysis

Using good planning and intuitive interfaces  
to bring GIS to the masses

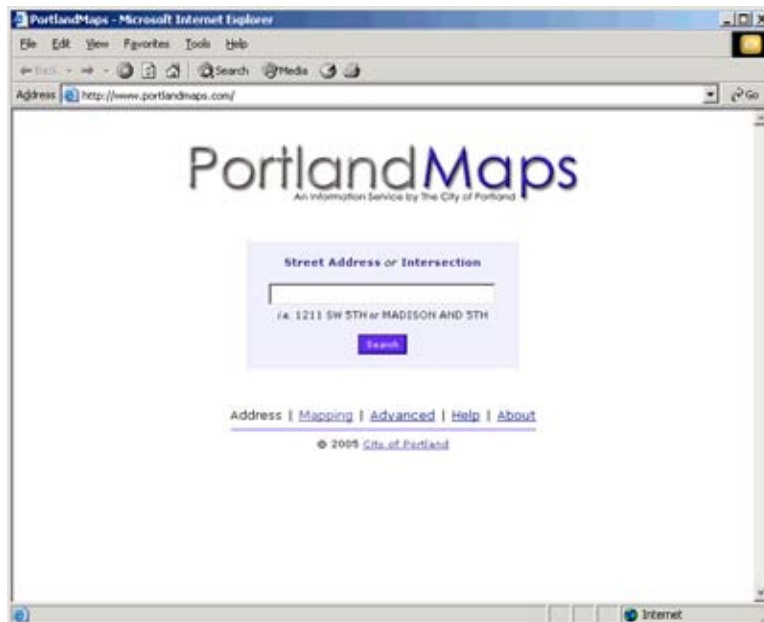
Phillip Holmstrand  
pholmstrand@ci.portland.or.us  
City of Portland, Oregon

## Introduction

4 years ago, a bunch of us from the GIS group at the City of Portland were sitting having lunch. We had just finished completing the "Internet Mapping" portal on the City of Portland website, a half dozen or so datasets each with a mapping interface. In each you could zoom and pan around the Portland area looking at the data, perform address searches, buffer tax lots, perform identifies on many layers at once, etc. Although we knew we had a great set of tools online, we weren't really sure that anyone outside of the GIS community would really "get it" enough to use it. We had spent years perfecting the interface, but it still wasn't right.

Then the idea came: build an interface that was extremely simple, with a start screen that simply asked you to enter in an address. Give it its own domain that's easy to remember, link in every dataset we have. Make it SIMPLE, and EASY. Our inspiration was the Google search engine, which (at the time) was just starting to blossom into a popular tool. To pay homage to our source of inspiration, we came up with the development name "Moogole" for our new service.

In less than 6 months we had a full production service, with assessor info, crime data, aerial photos, zoning, tax maps, hazard data, basically any and all data we had we put online. It was amazing how fast we could put together new pages for the service and link them in. Previously it had taken 2 years of painstaking work to perfect the previous "monolithic" or "map centric" interface, but with this new "report based" interface it was so easy.




PortlandMaps search screen was built to resemble Google in its simplicity

When we were done with our first version of the site which we released onto the internet as PortlandMaps.com, we knew we had something special. Without any fanfare, we linked it from our old website and waited. After a few months we took down our old map centric portal, even though it still contained features that the new one did not, and its funny, but nobody complained! PortlandMaps did have a bit less of the power user features of the old portal, but people didn't need all that. All they needed was the information and maps presented in an easy to use interface.

Over the last 4 years the service has grown, almost entirely by word of mouth, to a behemoth of Web GIS. Over 40,000 unique users per month, 160,000 visits per month, and 2 million page views per month. Ask people around Portland, they know PortlandMaps. Not only are most people familiar with it, but user polls have revealed that nearly 20% of our traffic comes from users use it as part of their everyday tasks. For city workers it has become an invaluable tool, we know just how valuable it is when we have some downtime and our phones and inboxes explode with questions "When will PortlandMaps be back up?! I can't do my job without it!"

Graph of PortlandMaps Visits over a 30 day period



Putting GIS on the web that was more than a novelty or a tool for power users was the real challenge with PortlandMaps. Most of our users don't know what GIS is, much less that they are using one every time they use PortlandMaps. Their main concern is to get the information they need and get out, and that's what we tried to do.

So where am I going with all this? Well in this paper I'd like to present what my 8 years of experience in Internet Mapping have taught me, and what recommendations I'd give to someone building a new Web GIS system, or taking an existing one and making it easier to use, easier to maintain, and more powerful.

## Implementing a Web GIS

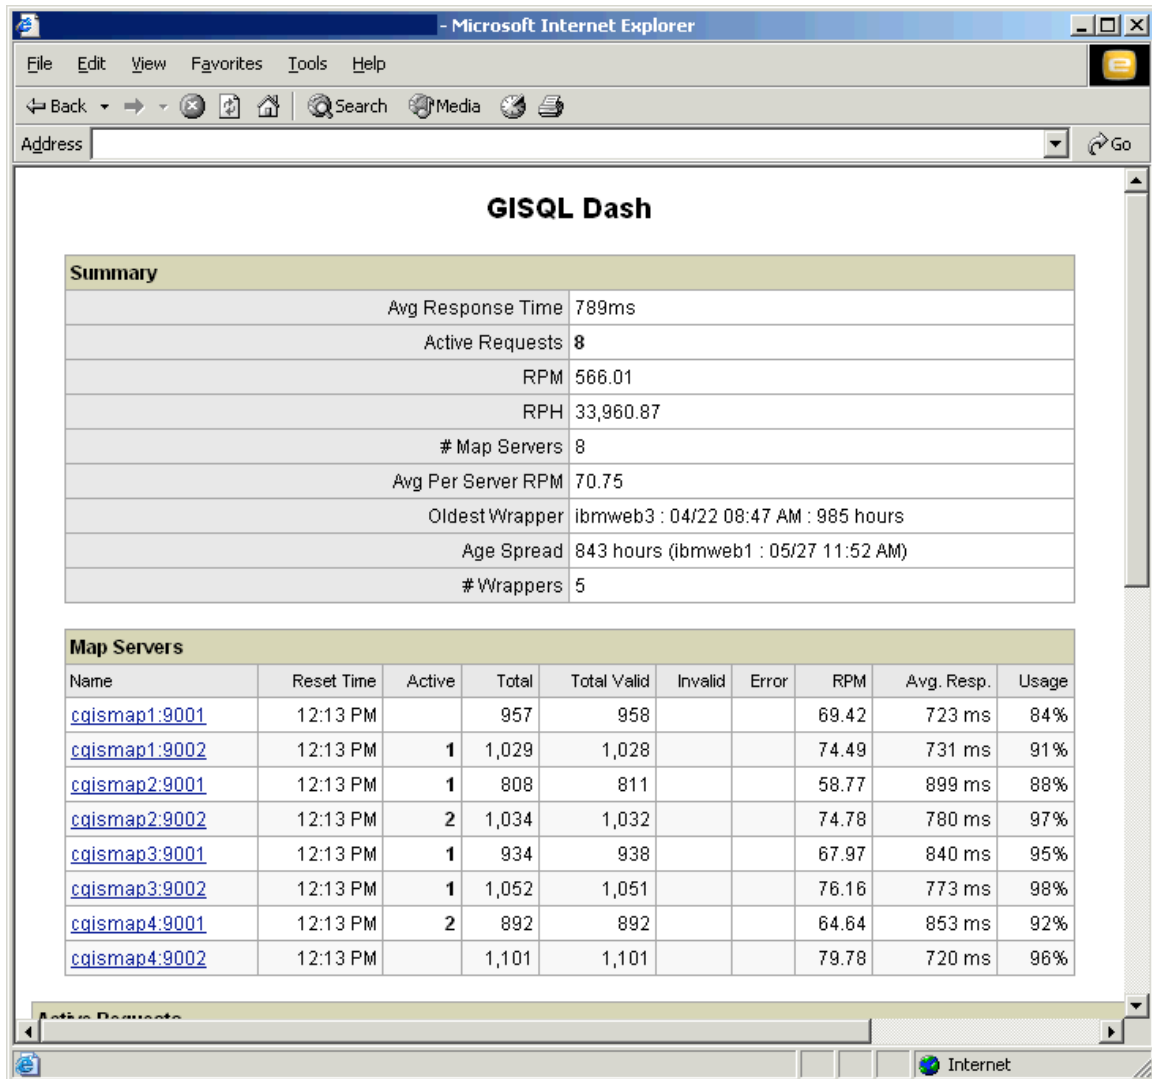
Putting your GIS on the Web (or even just your Intranet) can give anyone with a web browser access to powerful GIS analysis tools, but implementing a Web based GIS even when using powerful tools like ESRI ArcServer or ArcIMS is not a simple task.

First you must consider: What is the business need are you trying to provide the solution for? What data needs to be access? What types of conclusions will be drawn from your data? What types of automated analysis would help? Who are your users, are they experts or non-technical?

What methods will your users be using to search for data? If you start your application design based off of these questions, you will be off to a good start.

Second, take into account what type of performance and reliability is required. How many simultaneous users do you need to support? Here's a crude way of figuring out how many users you can support:

*\* If your map server and web server can generate one map/page per second, and users request a new map every 10 seconds on average, then your system will likely be able to support up to 10 simultaneous users. Of course your results may vary!*



We've had to develop a number of tools internally to manage requests to our map server cluster, allowing request level QA as well as enabling us to closely monitor traffic and performance real-time

The reliability requirements of your server are an important consideration too. What happens if your server goes down? is it sufficient to wait until someone can manually restart it? For most the answer will be no, so it's a good idea to look into some monitoring/alert software that can monitor your map server and run scripts to restart it if it goes down. We use Mercury SiteScope (<http://www.mercury.com>), it's an incredibly simple but powerful tool. I highly recommend it, our entire system would be nowhere without it!

Trust me, your map and web servers will go down, especially if your site has lots of traffic. Our map server processes restart as often as twice per day. Luckily our system has been built to handle it. Actually even if half our map servers goes down, our users will likely never notice anything happened. We have a middle tier "wrapper" that handles all requests and does QA on every one. If a particular map comes back blank or with an error, an alert is sent to SiteScope to restart the offending map server, and then the request is redirected to another map server. Of course a design like this is only possible if you choose to run many map servers in parallel.

The screenshot shows the SiteScope web interface for monitoring a group of servers. The main heading is "Monitors in the Group: Servers: Map Servers". Below this is a table of monitors with the following columns: State, Alerts, Reports, Status, Name, More, Edit, Refresh, Updated, and Del. The table lists 12 monitors, each with a green state indicator, a warning alert icon, and a report icon. The status column provides details on CPU and memory usage, and the name column lists the specific monitor type (e.g., CPU Utilization, Memory Utilization) for each server.

State	Alerts	Reports	Status	Name	More	Edit	Refresh	Updated	Del
●	⚠	□	48% avg, cpu1 48%, cpu2 49% (static baseline, 4.6875)	CGISMAP1 - CPU Utilization	Tools	Edit	Refresh	10:31 AM 6/2/05	✕
●	⚠	□	16% used, 10148MB free, 0.00 pages/sec (static baseline, 32.5083)	CGISMAP1 - Memory Utilization	Tools	Edit	Refresh	10:34 AM 6/2/05	✕
●	⚠	□	24% avg, cpu1 23%, cpu2 25% (static baseline, 4.6727)	CGISMAP2 - CPU Utilization	Tools	Edit	Refresh	10:36 AM 6/2/05	✕
●	⚠	□	26% used, 5530MB free, 0.00 pages/sec (static baseline, 32.4581)	CGISMAP2 - Memory Utilization	Tools	Edit	Refresh	10:31 AM 6/2/05	✕
●	⚠	□	20% avg, cpu1 19%, cpu2 21% (static baseline, 4.8379)	CGISMAP3 - CPU Utilization	Tools	Edit	Refresh	10:37 AM 6/2/05	✕
●	⚠	□	39% used, 3608MB free, 0.00 pages/sec (static baseline, 32.0306)	CGISMAP3 - Memory Utilization	Tools	Edit	Refresh	10:37 AM 6/2/05	✕
●	⚠	□	34% avg, cpu1 35%, cpu2 34% (static baseline, 4.7409)	CGISMAP4 - CPU Utilization	Tools	Edit	Refresh	10:33 AM 6/2/05	✕
●	⚠	□	34% used, 3953MB free, 0.00 pages/sec (static baseline, 31.5237)	CGISMAP4 - Memory Utilization	Tools	Edit	Refresh	10:34 AM 6/2/05	✕
●	⚠	□	2% avg, cpu1 2%, cpu2 3% (static baseline, 6.3919)	CGISIMS1 - CPU Utilization	Tools	Edit	Refresh	10:37 AM 6/2/05	✕
●	⚠	□	15% used, 2083MB free, 116.02 pages/sec (static baseline, 14.1117)	CGISIMS1 - Memory Utilization	Tools	Edit	Refresh	10:36 AM 6/2/05	✕
●	⚠	□	2% avg, cpu1 2%, cpu2 2% (static baseline, 6.5139)	CGISIMS2 - CPU Utilization	Tools	Edit	Refresh	10:36 AM 6/2/05	✕
●	⚠	□	18% used, 2000MB free, 171.39 pages/sec (static baseline, 17.244)	CGISIMS2 - Memory Utilization	Tools	Edit	Refresh	10:31 AM 6/2/05	✕

We use Mercury SiteScope for monitoring and automatically restarting (in case of failure) our map servers

Your third consideration in designing your Web GIS is what web application development environment you will use to build it. If you are using an ESRI based map server you will likely have SDK's available for Java, .NET, ASP, or ColdFusion. It is important to know that you can mix and match your web development language with the language your server is operating in. For example at the City we use a customized version of ArcServer written in C++ and VB, but our web development environment is in ColdFusion. Our middle tier layer is written in Java. The three components work together flawlessly because they all communicate over standard HTTP.

In picking your web development environment, don't get too caught up in the marketing of "most powerful, latest and greatest." Just choose based on your group's experience, and what looks like it will support everything you need to do. The web is a very primitive environment, and you really don't need as much power as you think you might in your development environment. Some languages like Java have been called the "SUV of programming tools": Capable of doing it all, but mostly just used to do simple tasks. Java servlets do a great job in our middle tier where we need a fast and powerful process that gives us a lots of flexibility, but for our presentation layer ColdFusion is much faster to develop and easier to maintain.

Now you have your business case, your performance and reliability requirements, and your development environment(s). You should have a good idea of what its going to take to put this all together. If you don't, here's a hint: Its not going to be easy.

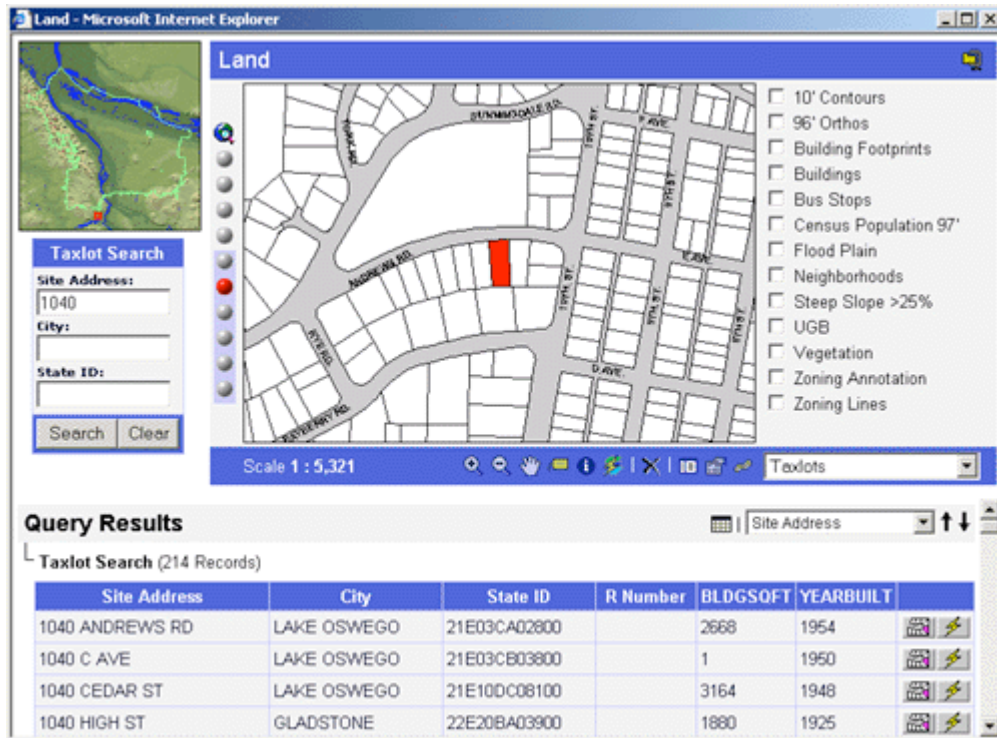
Many organizations may dream that they can buy a piece of software, install it, and get something workable right out of the box. Well for some this might be true, but usually these cases end up becoming more of a novelty map service than something real. To really get something useful out on the Web that will appeal to a large audience (and you should want to appeal to a large audience, otherwise why bother putting it on the Web) you need to put together an application that easy to use and helps people answer their questions.

## **Evolution of Web Based Mapping Interfaces**

When the first web sites with Internet Mapping based on ESRI technology debuted, they had relatively primitive interfaces. Simple HTML based forms without any JavaScript or DHTML to enable nifty features like dragging the map to pan. Some didn't even allow for any kind of direct map interaction, but rather used an address search function to center the map and zoom it on an area.

Then as browsers evolved the web mapping interfaces did as well, at least those focused primarily on GIS. (There were others with a broader appeal like MapQuest that kept relatively simple interfaces.) In 1999, before I joined the City of Portland, when I was still working at a software company, myself and a few others developed an interface that allowed for a desktop GIS like experience in a web browser. The interface resembled ArcView, with a map at the center, zoom buttons, on/off layer options, legend, "Map Tips", identify, geocoding, buffers, etc.. The interface took extensive use of the version 4 web browsers ability to refresh images without doing a full page reload, so using the interface seemed more like using a desktop app than browsing web pages. Since this type of interface is essentially built entirely around navigating a map, I will call this a "Monolithic Map" interface.

My colleagues and I were very proud of our creation, it took over 2 years to perfect, but when we were done we had something that we thought was very cool. Every GIS group we showed it to was impressed, and it ended up becoming a successful product that continues today. Shortly after I left the company and went to work for the City of Portland. When I arrived my first task was to adapt the product I used to be a developer for, to work for the City of Portland.

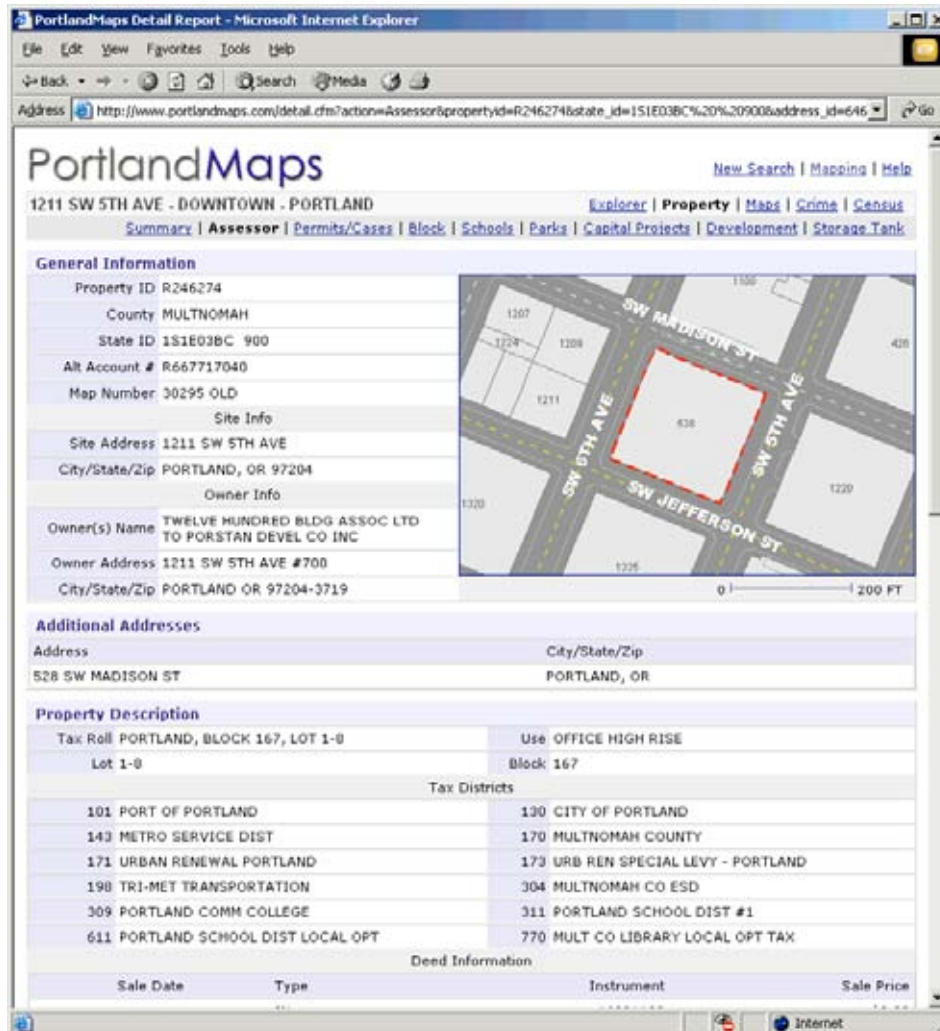


MapWorks, old City Web GIS portal, example of Monolithic Map interface

I had no problems quickly setting up the map server and customizing the interface to fit our needs. However after we had setup the service, we realized that although it worked well for GIS professionals, but made no sense to the average Internet user. The interface was far too complex and bulky. Performing a simple task like identifying property information took multiple steps and was not intuitive. Also because of the structure of the monolithic map interface, it didn't allow for easy customization or extending. Virtually every type of report or detail required popping up a new window, which did not link back to the calling interface.

The monolithic map interface, though it may be familiar to GIS professionals, was not right for our needs and the non-GIS user.

We decided to go back to the drawing board and re-think our approach. Our problem was how to more closely join the map interface and the detail screens that had all the information people needed. Navigation and information, joined at the hip, in an easy to use interface. The result is something we call a "Report Driven GIS." Maps and detailed database information are shown together, click on a map on a detail screen and it takes you to an interactive interface to change location. From that screen there is a link back to the detail screen. It's the bi-directional interaction between report and map interface that we wanted, and the real beauty of it is it could be developed using the same very simple HTML that we used 3 years previous on the very first internet mapping interfaces.



PortlandMaps: New City of Portland Web GIS portal, example of Report Driven interface

Because the interface was so simple, it took us only 6 months to have an entire site built using some 20 different datasets including Property, Photo, Zoning, Crime, Hazard, and Utilities. Over the years we've added more data like Building Permit information, Census data, and even complex reports like break downs of property information over a given block.

To easier understand what I am talking about when I say Report Driven GIS, go right now to <http://www.portlandmaps.com> ... enter in the address provided in the example, or another address in Portland if you have one. Browse around and you will see what I am talking about.

Done? Well okay, if you skipped it then let's just say that the interface we developed allows users to generate a report for a given coordinate in our GIS, and then interact with it. You can select from any one of our two dozen datasets by clicking the links at the top of the screen, much like you would navigate a website. It's intuitive because it uses the Web in the way it was meant to be used.

Changing locations on the map is done by using the "Explorer" page, which you can access by the link at the top of the screen or by clicking on any of the maps in the report. You can also go back to the homepage and enter in a new address or intersection. We have found that most people don't really browse maps too often, usually they are interested in a particular location,

usually a piece of property. Once they find it, they navigate through all the different datasets available. The Monolithic Map interface makes navigating the map easier than the Report Driven style, but if most people are just interested in a single location, then it makes more sense to base your interface around that.

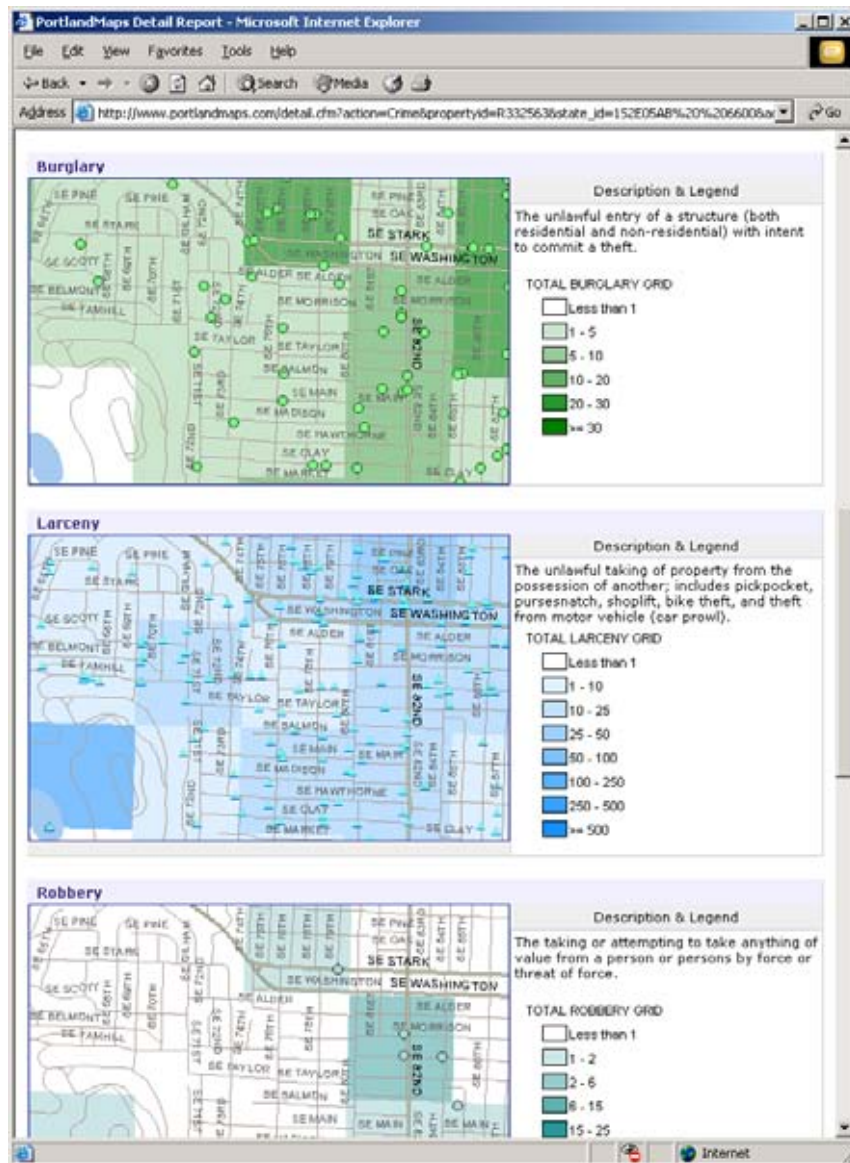
Try entering in an address that doesn't exist, like "7601 SE Yamhill" ... Since the address doesn't really exist, PortlandMaps will geocode it and you will be presented with a page that shows you all the addresses on the block that shares the range with the address you entered. We could tell the user that their address has been geocoded, but the only way they would know what "geocoding" means is if they were a GIS professional, and that's not likely! So instead we put them on a page getting them as close to the address they entered as possible and show them the addresses on that block. Maybe they mistyped the address, or maybe our data is wrong. Either way we are trying to get them as close to what they asked for as possible.



Aerial Photos are indexed by year and resolution, streets and tax lots can be toggled on/off easily



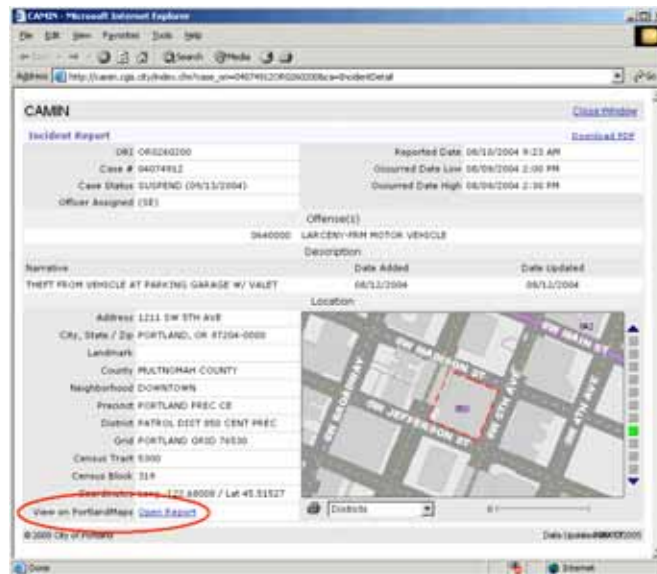
One feature that our Report Driven GIS interface gave is the ability to simultaneously view many maps showing different datasets. This is a very powerful tool when comparing and contrasting different data, without the complexity of having the user manage the turning on/off groups of different layers. Of course with as many as a dozen maps per screen on PortlandMaps, the backend needs to have the server power to support it. Currently we run 4 dual processor 3.2 GHz map server machines, with 2 map server processes per machine (we run two processes to take better advantage of the multiple processors.) Even with all of that, we are looking at adding two more machines as we are almost near capacity with our existing cluster.



**Multi-Map Report:** Possibly the biggest advantage of the Report Driven GIS is the ability to simultaneously show many maps of the same location but with different datasets

We've taken the Report Driven GIS concept to all of our Web GIS applications, over 3 dozen in all, most linked in through PortlandMaps. Since PortlandMaps links in all of the different datasets, you get a great synergy every time you add more. For example Portland Development Commission had created frontage photos (snapshot digital photos of the fronts of buildings) for all of downtown. We linked this in to PortlandMaps, so it instantly became available to every other

application linked to PortlandMaps. An example: When the Police Bureau uses the Crime Analysis tools we built them to look at a bank robbery, they can click the "View in PortlandMaps" link and be presented with an actual photo of the front of the bank!



Report Driven GIS allows easy linking of many datasets and separate applications together, such as linking crime incident information from the Police Bureau in CAMIN to a photo of the building taken by the Development Commission in PortlandMaps

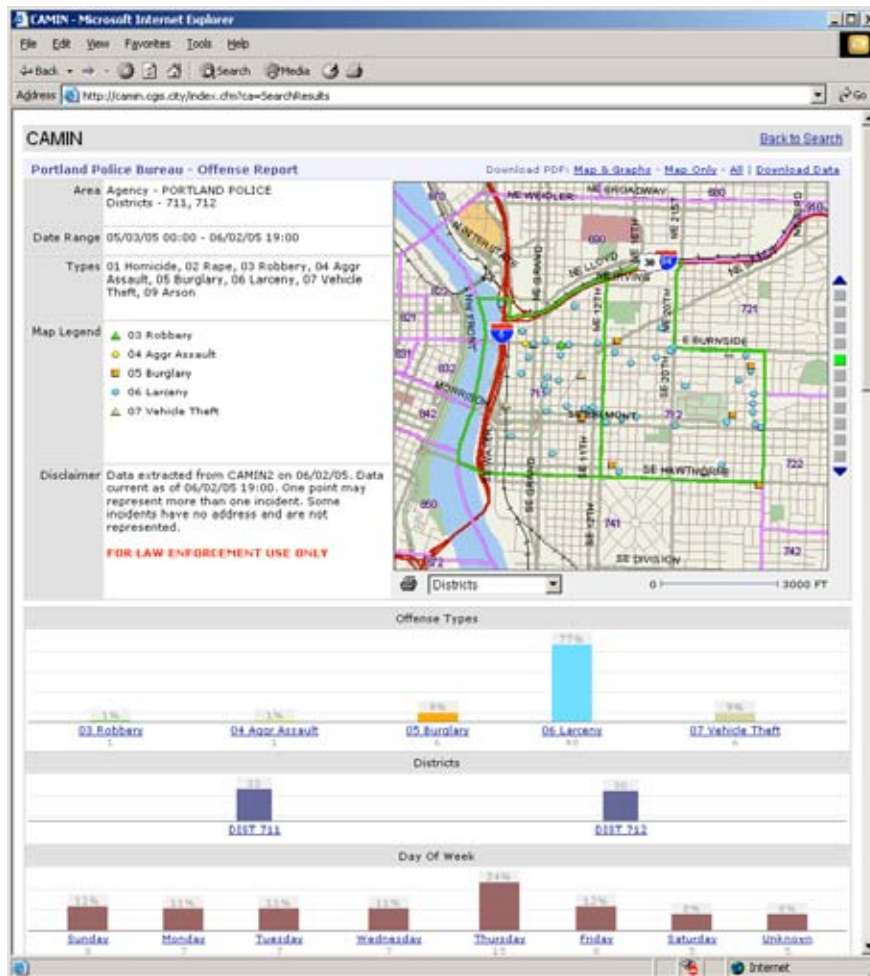
The real promise of GIS is helping people draw conclusions by linking many sets of information together. The Report Driven GIS interface lets you do this without all the hassle of cramming it all into an interface around a giant map. Instead, the users navigate through a simple interface that resembles the rest of the web, to see the different datasets and find common trends between

them. Some contain maps, some don't... more to the point they all contain the necessary amount of information for the user to solve whatever question they might have.

## CAMIN and Other Web GIS Projects @ City of Portland

Since PortlandMaps was so successful, we have taken the same concept to all of our web GIS projects. Most of the time when people here at the City come to us with an idea for a new Internet Mapping application, it simply becomes another page on PortlandMaps. They are happy because their application instantly inherits all the power within PortlandMaps, and we are happy because we don't have to do much work!

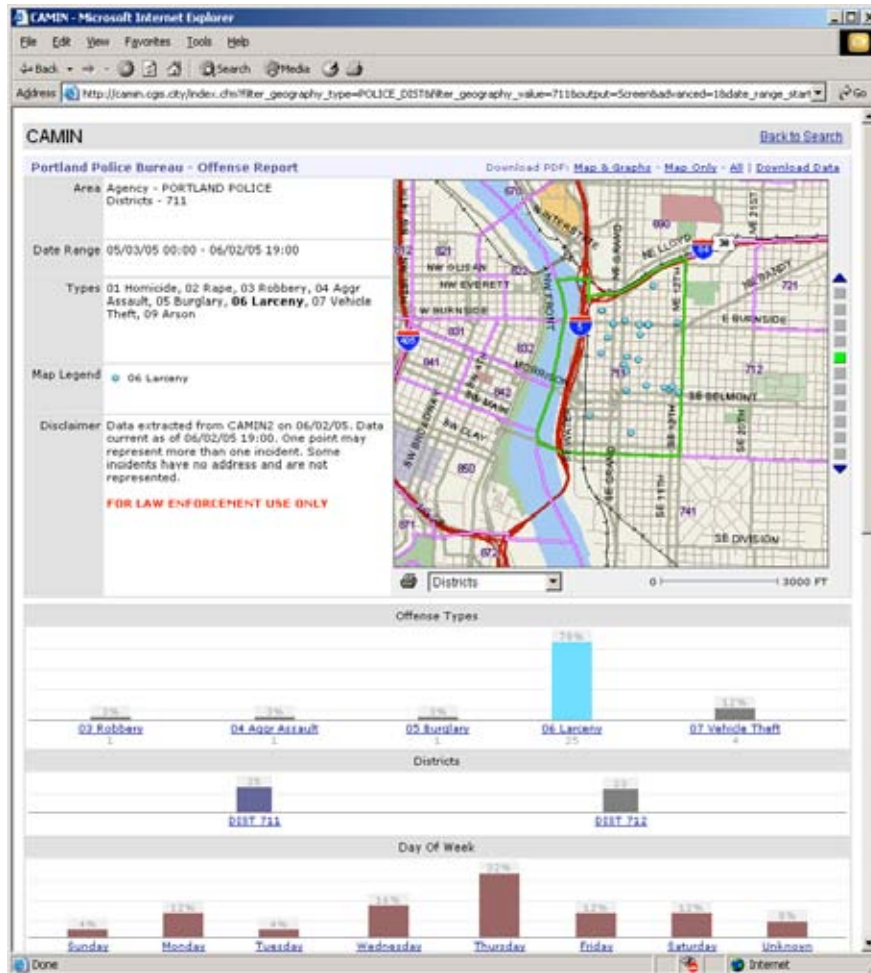
Our latest application where we used the report based concept, was for the City's new internal crime analysis system: CAMIN. CAMIN, which stands for Crime Analysis Mapping Information Network, is actually its own tool separate from PortlandMaps.



CAMIN mixes a crime analysis report, with a dynamic map and graphs which allow the user to carefully sort and filter the data

With CAMIN we were presented the challenge of exposing a huge amount of data (the entire Portland Police Bureau crime database) to a as wide an audience within the Police Bureau as possible. The goal was to get any police officer to hit the ground running, and in just a couple hours of training to be proficient. We also had to support an "advanced" user class of crime analysts that wanted the tool to support any type of question they could ask it, like: "How many

car thefts have their been within 500 feet of 5th and Madison in the last 10 years between 10pm and 2am?"




In the scenario above, the user has taken the original dataset which included multiple crime types and two different districts, and filtered down to just Larceny and district 711 by clicking on the graphs, the graphs and map update to reflect the new filtered dataset

Since most of the interface of CAMIN is driven by the report screen, no additional development had to be done to generate a "printable report" version of the data. Additionally we added the ability to download a PDF version of the report so Police Officers could save it off for reference, or use the enhanced PDF print environment to more closely tweak print settings.

## Portland Police Bureau - OFFENSE MAP

0 1500 FT



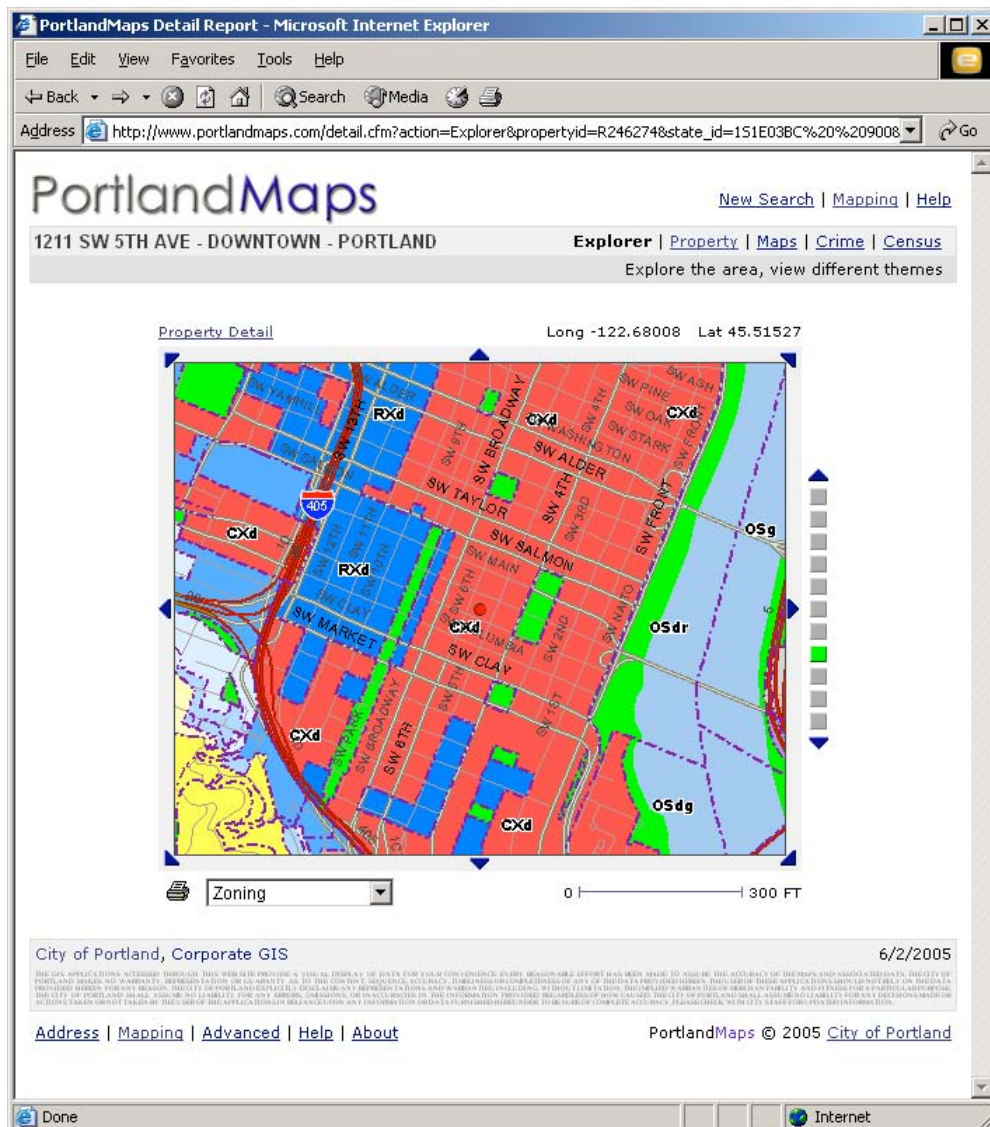
With ArcServer/ArcObjects you can get some really nice looking maps

## Reuse of Components

Another major advantage of the Report Driven GIS interface is how easy it is to develop reusable components. We were able to build CAMIN fairly quickly by re-using many of the components that already existed for PortlandMaps. You can build a Monolithic interface using reusable components, but because of all of the frames, JavaScript, and DHTML flying around, you will find that things get complicated very quickly. Because the Report Driven interface is based on a very simple design, components too are very simple.

We have two good examples of the same components in both CAMIN and PortlandMaps. The first is our address search tool. The address search tool is responsible for parsing a string that could be an address/intersection, breaking it down into its individual components, and searching the address database and/or geocoding it to find the address on the map. If you are familiar with developing GIS applications, you have probably done something like this in the past, likely many times over. We chose not to re-invent the wheel every time, so we simply wrote the component to be as flexible as possible, and re-use it every time we need it.

The second major re-useable component we put together is the map control interface; we call it the "Explorer." The Explorer is what you use to pan, zoom, and change themes on an interactive map. We built it so that simply by passing it a list of parameters you can control the data that is shown, extents of the map, and size of the map. The Explorer interface is probably the most complicated code in our entire Web GIS, and we didn't want to have to manage lots of different versions of it across our many applications, thus it made sense to build it as a re-usable component.



The "Explorer" interface shown here in PortlandMaps shares the same code with the one used for CAMIN

## Conclusion

Although we as GIS professionals may want a powerful map interface, most internet users have never heard of GIS. The key to unlocking GIS to the masses is making it so users can use a GIS without ever having to know what geocoding, layers, or features even are. Also because most information users are interested in is actually contained within a database, it makes more sense to base the interface around the database and let the map become an feature of the data.

I think that with a move to Report Driven GIS on the web, we can finally let GIS get out of the back-office and into the world.

## Summary

### Web GIS: Is it right for my organization?

- Considerations Web vs. Desktop:
  - Programming Required for Web GIS
  - Web GIS is not end-user manageable
  - Will it be useful? Or just a novelty?
  - Powerful Server(s) vs. Powerful Clients, where do you have the hardware now?
- Advantages over Desktop:
  - Easier to deploy to large numbers of users
  - Easier to update code
  - Better for older client computers or slower network
  - Easier to integrate with many data sources
- Disadvantages over Desktop
  - Interface not as flexible
  - Requires more server hardware
  - More monitoring tools needed

### Web GIS: Checklist for Success

- Use Case Questions
  - What questions will the application attempt to answer?
  - What parameters will the user use to find what they are looking for?
  - What output will the user expect from the system?
  - Are there different classes of users in the system? Expert and Novice?
  - What options should the system provide for different user classes?
- Ease of Maintenance, Performance and Reliability
  - How many simultaneous users does the system need to support?
  - How fast can maps and web pages be processed?
  - What hardware/clustering is necessary to achieve the performance goals?
  - What monitoring/alerting software will be used to achieve the desired uptime?
  - Will it be necessary to stop servers while new data is loaded?
  - How often will data need to be updated?
  - Will schema of data possibly change when it is updated?
- Development Environments and Mapping Platforms
  - What Map server package will be used? ArcIMS? ArcServer?
  - What SDK's are available for the given map server?
  - What web development language(s) will be used?

- Remember to focus especially on:
  - Ease of use to end user
  - Reusability of code, and flexibility of expansion
  - Ease of maintenance

### Web GIS Interface: Monolithic Map or Report

- Monolithic pros:
  - Similar interface to desktop GIS
  - Allows user to more closely "tweak" map
  - Can sometimes be provided "out of the box" by the map server software
- Monolithic cons:
  - Not very "web-like"
  - Not intuitive to non-GIS users
  - Often clumsy, buggy, especially in older browsers
- Report pros:
  - Very easy to integrate with many datasets
  - More "web-like"
  - Intuitive to non-GIS users
  - Allows for easy printing
  - Works well in all browsers, requires less DHTML wizardry
  - More suitable for displaying database data
  - Compatible with non-desktop browsers like PDA's or cell phones
- Report cons:
  - Requires more server power (if multiple maps per page are used)
  - Requires custom interface development
  - Not as easy to "tweak" map
- Monolithic vs. Report Conclusions
  - Report interface gives a more web like user experience, with easier tie-ins to many datasets.
  - Monolithic interface may be familiar to GIS users, but can be overwhelming to novice users.
  - Monolithic interface requires extensive JavaScript and DHTML code to work well, although pre-built packages are available, they still need to be customized.
  - Monolithic interfaces break down when it comes to dealing with large selection sets easily, they get very complicated very fast.
  - Report based interfaces will likely require more server end power on both the web and map server, especially if multiple maps per page are used.
  - Monolithic map interfaces do not allow for complex search/results steps.
  - Report Driven GIS is more like using a web site, and thus more intuitive to web users.

### Technology Used

- Server Software
  - ESRI: ArcServer, ArcObjects, ArcSDE
  - Microsoft SQL Server 2000
  - ColdFusion MX 7
  - Java Tomcat
  - Microsoft IIS 6
  - Microsoft Windows 2003



- Server Hardware
  - Map Servers: 4 x Dual 3.2ghz IBM Blades
  - Web Servers: 4 x Dual 3.2ghz IBM Blades
  - Database: 2 x Eight Way 2.8ghz IBM

## **Acknowledgements**

### **Map Server Development**

Reggie Wilbanks – GeoNorth Inc.

Kellie Hauger – City of Portland, Corporate GIS

Mike Quetel – City of Portland, Corporate GIS

### **Map Server Middle Tier**

Anthony Dotson – City of Portland, Corporate GIS

### **Data Management**

Mitch Vanderperren – City of Portland, Corporate GIS

### **CAMIN Project Team**

Mark Baird – City of Portland, Police Bureau

Steve Beedle – City of Portland, Police Bureau

**Also thanks to:** Rick Schulte, Marshall Payne, Dan Bauer, Herts Chen, and Aaron Kilbey.

### **About the Author**

Phillip Holmstrand is Lead Web Developer for the City of Portland. In 1997 while working at GeoNorth Inc, he worked on the team that put some of the first GIS mapping tools on the internet. He was lead interface developer on the GeoNorth product MapOptix. In 2000 he joined the City of Portland, Corporate GIS team. While at the City he's headed a number of large scale web applications such as the City's CMS PortlandOnline.com, and the City's GIS portal PortlandMaps.com. Phillip lives in downtown Portland and enjoys Kiteboarding, Wakeboarding, Snowboarding, listening to music, and playing Xbox.