

t1.8

Confessions of a GIS-enabled B2B web site

Adrian Strahan, Business Development Manager, Sitescope Ltd

Introduction

Web-based services offer an enormous potential for the provision of information. More and more of these services are location-based, where the information is specific to a particular geographic location.

However, a web site's success is measured by how quick it is to load and how relevant the information it provides is. Not to mention how easy it is to use and how reliable it is!

On a web site offering a free service, one can be forgiven if things sometimes go wrong. But when trying to offer a service for a fee, these issues become very important and one cannot afford things to go wrong. The user becomes a customer and the customer demands a professional service that is quick, reliable, instils confidence and makes them want to use the service again and again.

In terms of usability, good web design helps to a certain extent but the addition of maps and GIS allows the information provided to become much more meaningful and relevant to the user. This is particularly true when applied to the property market. However, we cannot afford to provide a slow service – reliability and speed is what all customers are after and if they don't get it they would justifiably go elsewhere. So the capabilities of any GIS used have to be good enough to meet these requirements.

In October 2000, Sitescope launched "Homecheck Professional", a business-to-business (B2B) web site (www.homecheckpro.co.uk) that provides on-line environmental reports, aimed at conveyancers and property professionals dealing with both commercial and residential property transactions. Since the launch, Sitescope have become the largest provider of on-line environmental reports in the country. But to achieve this feat, we have had to utilise the most effective technology in the most effective ways.

This paper therefore describes some of the issues we have dealt with to provide such a service and, in particular, it describes how our GIS system has helped us to meet the demand we currently have today.

Know Your Customer

Our users are conveyancers, solicitors and property professionals who use the Homecheck Professional service to help them identify any potential environmental issues that may affect their clients' present or future property.

Interestingly, our own research has shown that the people who actually use the web site tend to differ between different types of practitioner. Within conveyancing companies for example, the conveyancers themselves tend to be the main users of our services, whereas solicitors are more inclined to ask their secretaries to produce the reports.

However, practically all of our users are familiar with maps since they have to deal with site plans or land registry documents on a day-to-day basis. They usually also have local knowledge of the areas they're dealing with – a luxury our national web site cannot hope to have.

So our users know their way around – but they don't necessarily have, nor necessarily want to have, any GIS experience or skill. So how can we get them to precisely specify the location for their report - in the simplest, most user-friendly way possible?

Site Identification

Site location is a crucial element of any property information site – it's a task that most people familiar to GIS will be quite happy with. But the World Wide Web is a challenging environment for application developers and a web site can potentially be used by thousands of users – users who have varying levels of knowledge and skills. So when a web site - where there's no human contact at all - asks a “spatially-challenged” individual to specify a location, the task could be very difficult.

Location, Location, Location?

In order to perform a search about a property or site, one has first to identify where that property or site is. More specifically, and for our purposes, one has to identify a grid-reference within the property or site. If the user already knows this, then all's well and good. But very few users have this information, simply because they don't have access to it to an accurate enough level. And even if they did, they may not understand what the grid reference actually was or meant.

So, in the majority of cases, we have to offer the user a more human-friendly way of specifying a location. An address usually suffices, but to accurately locate the address, we generally need its postcode and the name or number of the house, or the name of the organisation. In a lot of cases, this information isn't always available. This may be because the user only has part of the address – without the postcode for example; they may only have the name of a place (such as a locality); or they may only have a plan of a plot of land that hasn't even got an address yet.

Occasionally, the user only has an idea of the postcode – maybe just to district level, e.g. “SE25”. And in the worst case, the user may have a postcode and the property's name or number, but either or both of these aren't recognised by the system (maybe because the latest Address-Point update hasn't been received or loaded yet).

Remember, if we were providing a telephone service for example, where human interaction is used, we could probably have a list of questions we could ask the user to help identify the location required – helping to slowly narrow down the options until the correct location is found. But on a web site, we have no such luxury. So we have to employ a number of different methods to cover all of the eventualities. It's the old “80:20 rule” – 80% of the work is required to cover 20% of the situations.

I could write a paper on just this one subject, so here's just a brief outline of what we do for each situation:

Postcode and Property Name/Number Known

Let's start with the easiest situation first. Using Address-Point from the Ordnance Survey (OS), one can identify a grid reference of the required site to within 10 centimetres when given a postcode and the property name or number, or organisation name.

Address Known, But Not the Postcode

Having an address without the postcode is a common situation. The idea of a postcode, although meaningful to the Royal Mail, isn't very user-friendly, so some postcodes are not known, or worse still, incorrect. It's also quite common for people to misinterpret the characters – specifying the letter “O” instead of the number zero for example.

By searching Address-Point for the street and town, one can give the user a list of matching entries with their postcodes. When they select one of these, a list of property numbers and organisation names can be displayed. Further selection of one of these will give us the grid reference we require.

Only the Name of The Place is Known

Generally, this occurs in the more rural areas where the property could be, for example, that of a farm or group of cottages. Occasionally, it must be used when the postcode or address is too new to be in our Address-Point database.

We use the Landranger Gazetteer from the OS to help us here. This allows the user to specify an area or feature that can be seen on the corresponding Landranger map. This gives us a rough grid reference, from which the user can fine-tune the required property's location.

Mapping It

So, where does the mapping come in? Well, until we have a grid reference, we can't really show any meaningful map. There really is no point in showing a map of the whole country and expecting the user to "zoom in" until they reach the appropriate level of detail. This may look slick but on a web site it's just a waste of time. If all the user really knows is a general location, they can enter the name of the nearest town or village instead and have it matched against the Landranger Gazetteer.

Once a grid reference has been identified a map can be shown – at an initial scale dependent on which of the above methods the grid reference was obtained from. For example, if the grid reference was entered directly by the user, or came from a match against Address-Point, we can assume that the location is accurate and that the largest scale mapping should be shown. However, if only a place name has been entered and matched against the Landranger Gazetteer, the first map to be shown should be the corresponding Landranger map (which, incidentally, will have the name of the place already displayed on it).

Our web site's map page allows the user to view maps between a number of set scales – from 1:3,000 to 1:500,000 – each showing the most suitable mapping that we have for that scale. We only offer set scales because there really is no point in allowing the user to say, zoom in to a scale of 1:1, or to zoom out to a scale of 1:10 million.

At the largest scale we show Internet Street Mapping (ISM) from the OS (see Figure 1). This allows the user to see individual building outlines – a very useful feature when trying to identify individual properties.




Figure 1: ISM map showing location of site

The Landranger map (Figure 2) currently caters for scales between 1:10,000 and 1:50,000 and helps show a general view of a location.




Figure 2: Landranger map showing location of site

The OS's 1:250,000 Travelmaster map (Figure 3) is used for scales between 1:100,000 and 1:500,000.

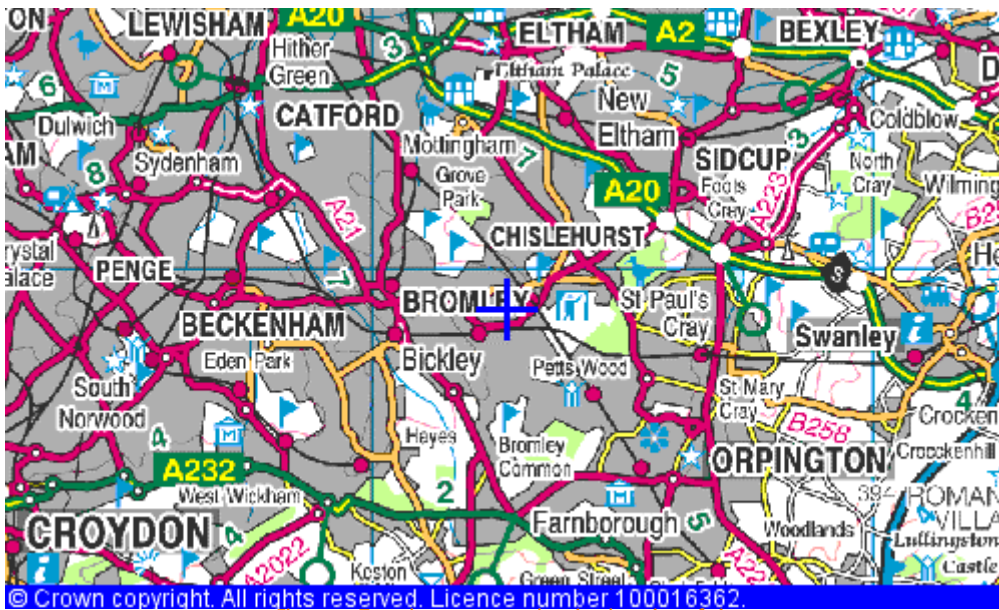


Figure 3: Travelmaster map showing location of site

Simply Does It

Note that the map page is very simple in its concept. There is no need for complicated JavaScript code to try and identify where the cursor is pointing or any need for polygon digitisation routines or vector download. There is also no requirement for a plug-in or applet. By using the minimum amount of functionality and coupling it together with a simple raster image that can be refreshed very rapidly, we are able to support many simultaneous users.

X Marks the Spot

As you can see from the examples above, when the map is displayed, we immediately centre it on the grid reference we have and place a cross at this location. If the location was obtained via the gazetteer, we ask the user to confirm the location of the property or site by clicking on the map at the correct location. The user is allowed to pan the maps and change viewing scales until the desired site has been located.

Even if the location was originally obtained from Address-Point, we still allow the user to move it. This is because they may not always agree with the location returned – and who are we to argue? The customer is always right!

The final location specified by the user will be the point from which the search will be made. This need not be the site of a particular building – it could, for example, be for a plot of land next to an existing address.

It's worth underlining this point. Although we may have used a grid reference, an address or the name of an area, to initially identify where the property or site is, the final location chosen by the user for the search may be different.

User Confidence

The map page is, I believe, the most essential aspect of the site location process. It not only allows us to ensure that the location required is correct; it allows the *user* to ensure that the location is correct. Remember that the user is about to spend some money and the last thing they want to do is request a report for the wrong place.

The Report

The report is generated in real-time once all the required details have been entered and confirmed by the user. As soon as it has been produced, the user is immediately shown a summary of the results of the search. A PDF version of the report is also generated automatically at the same time – a printed copy of this is what the user will most likely pass on to their client. PDF is becoming the standard way of distributing documents over the web. “Acrobat Reader” is a free application supplied by Adobe (the inventors of PDF), which allows one to view and print PDF files.

The Report Summary

The report summary highlights any potential areas of concern and allows the user to view the details of anything found within 500 metres (for a residential property search) or 1000 metres (for a commercial property search).

However, just showing a list of features is not very useful – so we use dynamically generated maps to show the locations of each different type of feature in relation to the search location and label each of them with a unique, sequential number (see Figure 4). Again, the location of the search site is shown as a blue cross, but this time, we also include the relevant search buffer.

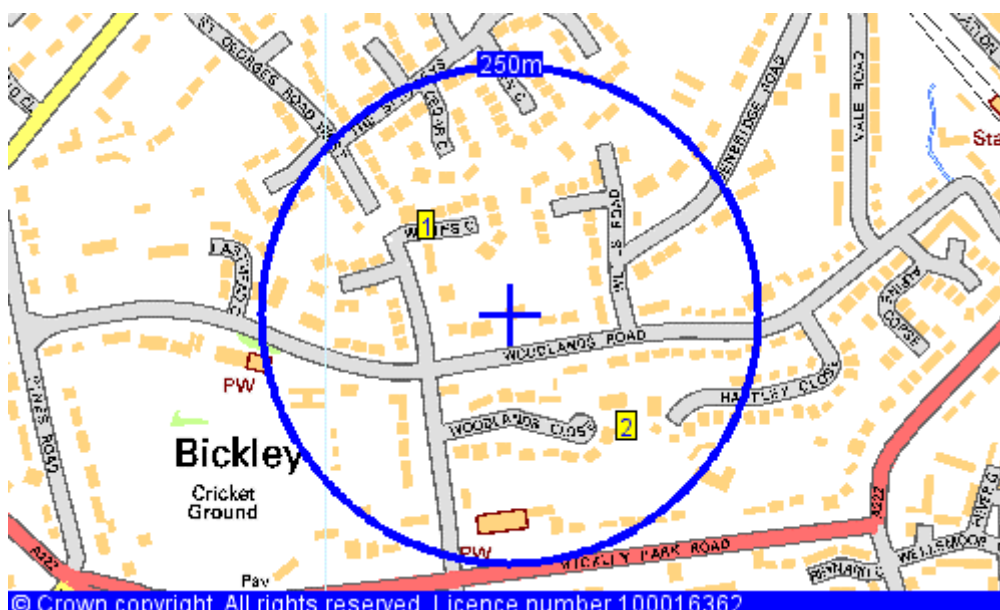


Figure 4: Map showing location of possible causes for concern within 250m of the search site

Each number corresponds to an individual feature in the search results list to allow the user to quickly see which are most likely (if any) to affect the search site.

The PDF Report

This contains similar information to that shown on the web site. All the details of everything that was found within the search buffers are shown, together with the maps showing the locations of the features where necessary. Additionally, the report shows a map indicating the full extent of the area searched (see Figure 5).

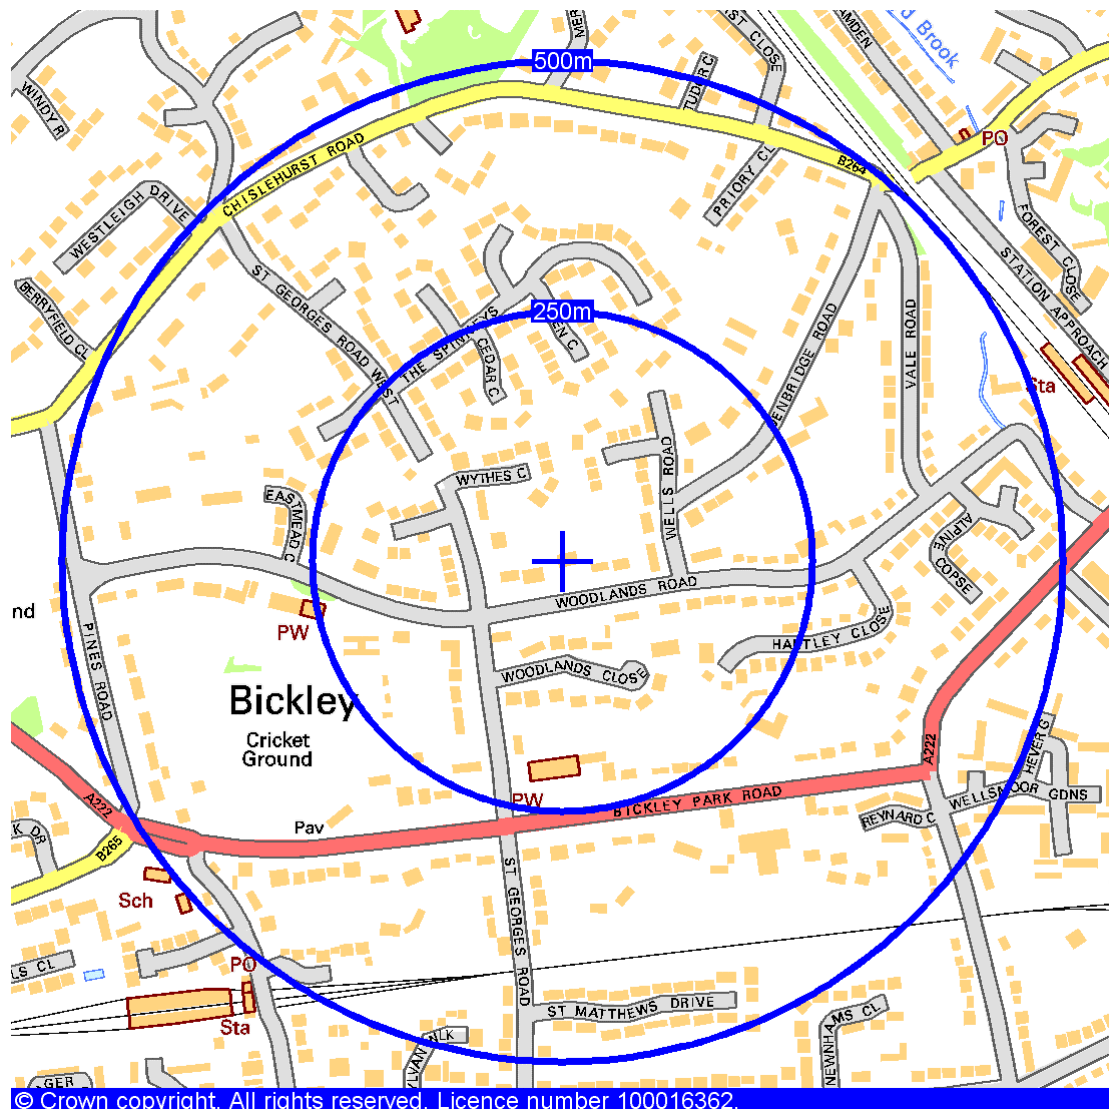


Figure 5: ISM map with search buffers

It should be pointed out that we use exactly the same mechanism as the web site to generate the maps for the PDF report. There is no special print engine routine or layout tool used - simplifying the process enormously.

Print Quality

Since the ultimate intention is to print the PDF report, the mapping has to be of a much higher quality than that used on the web site. In practice, a browser displays the maps (displayed as GIF files) at 96 dots per inch (dpi) and most printers today are capable of at least 300dpi. If we were to show an image from the web site in the report, the quality would be very poor.

So, in order to get higher quality, the images requested from the GIS are larger than normal – sometimes up to four times as big. When these images are inserted into the PDF report, they are automatically “shrunk-to-fit”. In other words, the PDF writing process manages the interpolation of the extra pixels. When the image is ultimately printed, the extra information has not been lost and the image quality produced by the printer is much higher than from a print of a normal web image.

For example, on the web site, we can really only view ISM maps up to a scale of 1:3,500 before the quality of the image becomes too poor to read the street names.

This is mainly because the source raster has not been anti-aliased (a problem we don't have with Landranger for example). In the PDF reports, however, we can show ISM maps up to scales of 1:15,000 simply by increasing the size of the image and showing this extra information within the same print area.

Whilst this ultimately increases the size of the PDF files, the quality of the final product more than justifies the means.

Area Features

So far, I've only discussed the display of point information on our maps. However, much of the environmental information we deal with is polygonal-based. This is where a map can really help the user to understand not only where potential risks lie, but also their extents. Figure 6 shows an example of some flood plain data.

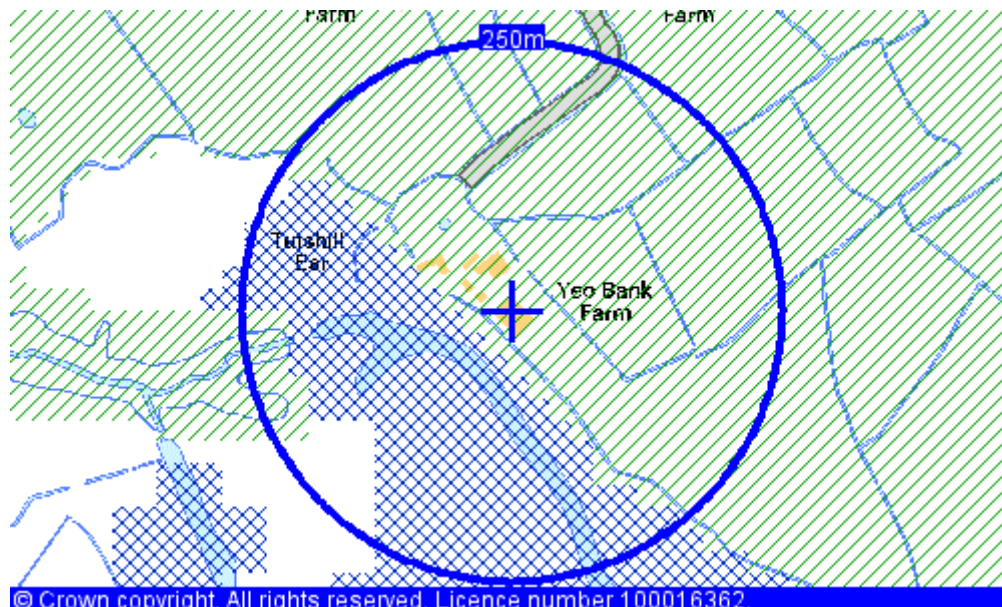


Figure 6: ISM map showing fluvial and tidal flood plain

The Future

Currently, only a co-ordinate is required for us to perform a search, and the functionality required to do this is very simple. But the future may require us to digitise a polygon indicating the boundary of the property or site in question - particularly if the site is much larger than the average house. For this to work, however, much greater thought must be put into the design of the interface to allow any un-trained individual (not just those with GIS experience) to select and digitise such features.

Also, there is the issue of having a suitable map from which to digitise a boundary from. At present, the only viable map that would allow this is Land-Line from the OS. However, this is an enormous dataset (especially when dealing with the entire national coverage) that requires a truly scaleable GIS that can also handle multi-user *Internet* deployment. At the moment, the jury's out on that one!

Conclusion

If you offer something for free on the Internet, people will forgive you when something fails. When you're charging a fee for a service however, customers expect a lot more and are less forgiving.

Customers will only pay for something if they think they are getting a good service. A B2B service over the Internet is no exception to this rule. It has to be robust and you cannot allow anything to fail – and that includes the GIS.

Mapping is an essential part of our on-line services. It enhances the user's experience and gives them confidence in the site location process. Within both our on-line summary and PDF reports, the maps also

help the user to immediately see if there are any features that may have a potential environmental impact on the search site, and show the positions of these features relative to it.

Delivering mapping over the Internet and within PDF reports in order to support and enhance our B2B web site is a very challenging task. We have had to be smart and have not tried to be too “flash” or unrealistic. The mechanisms employed within the site location process, for example, are simple and do not rely on the user having any digitisation skills. The technology behind the mapping is very simple, and does not rely on plug-ins or complicated code. The same process is used for both the on-line maps and those within the PDF report. This makes the system simple, efficient, fast and, above all, robust - handling thousands of map requests every hour – without fail.

If only all systems you see on the web today were like that.