The Dawn of the Location Enabled Web

July 6, 2009

Policy Posts are in-depth analyses on current tech policy issues from CDT experts. Sign up to receive the latest Policy Posts here:

---

1) Location Privacy

The ubiquity of increasingly high-powered mobile devices has already spawned the Internet's first generation of location-based services and applications. As the accuracy of location data improves and the expense of calculating and obtaining it declines, location may well come to pervade the online experience. While the increasing availability of location information paves the way for exciting new applications and services, the increasingly easy availability of location information raises several different kinds of privacy concerns. Ensuring that location information is transmitted and accessed in a privacy-protective way is essential to the future success of location-based applications and services.

Because individuals often carry their mobile devices with them, location data may be collected everywhere and at any time, often without user interaction, and it may potentially describe both what a person is doing and where he or she is doing it. For example, triangulation of an individual's mobile phone can reveal the fact that he was at a particular medical clinic at a particular time. The ubiquity of location information may also increase the risks of stalking and domestic violence if perpetrators are able to use (or abuse) location-based services to gain access to location information about their victims.

Location information can also be highly identifiable, even when it isn't directly associated with other personal information. For many people, there is one location where they spend their daytime hours (at work) and one location where they spend their nighttime hours (at home). After a day or two of collecting just those two data points about a person, it becomes fairly obvious whom those data points describe.

Furthermore, location information is and will continue to be of particular interest to governments and law enforcers around the world. Standards for government access to location information held by companies are unclear at best and far too low at worst. The existence of detailed records of individuals' movements should not automatically facilitate the ability for governments to track their citizens, but in many cases, laws dictating what government agents must do to obtain location data have not kept pace with technological evolution.

Testimony of Leslie Harris [4] (regarding DPI but containing a section about location) (April 2009)

Digital Search & Seizure Report [5] (February 2006)

---

2) The Dawn of the Location-Enabled Web

Apple recently announced the release of the iPhone 3.0 software, which is a free update available to iPhone users containing a number of new software features. With the release of the software, the latest version of the Safari web browser running on the iPhone will be location-enabled. This means that any Web site can ask Safari for the user's location, and Safari can provide it by using the

location positioning technologies built into the phone (including GPS, among others). Apple has implemented a simple interface (based on a draft of a W3C standard) that Web sites can use to request location.

Even before browsers started to become location-aware, Web sites have for years been using reverse-IP address lookups to obtain the approximate locations (at about city-level precision) of Web users. But with 40 million iPhone users, Apple'™s foray into geolocation marks the true beginning of an era when pinpointing many Internet users on a map - with the precision of a few meters, not a few miles - goes from complicated and onerous to simple and fast. This certainly will not work for all Internet users, but 40 million is a significant start.

This new development has some obvious privacy implications. CDT believes that location information should only be used on individual Internet users' own terms. Individuals should get to decide with whom they share their location, what that information is used for, whether or not it gets shared, and how long it's retained. Location-enabled technologies - including Web browsers - should be designed with privacy in mind from the beginning and with built-in user controls to allow individuals to manage their location data as it is collected. CDT has been working for years to incorporate some of these concepts into technical standards, originally in the IETF's Geopriv working group and more recently within the W3C Geolocation working group, which created the draft standard that Apple and other browser vendors are starting to use.

Although the initial attempts from Apple and others are highly protective of privacy in some ways, there is still much room for improvement in providing user control. With Safari on iPhone, each Web site that wants to use your location has to first obtain the user's permission not once, but twice. Those permissions are also reset every 24 hours. As far as consent goes, this is a really strong baseline.

But in terms of providing more granular control and transparency, the iPhone is lacking. There is no way for a user to see with which sites (or applications, for that matter) he or she has shared location. If a user visits a site and declines to provide location to it, the site may continue to prompt the user to provide location on every visit. It would be helpful for users to be able to have a whitelist of trusted sites that can always obtain the user's location, and a blacklist of untrusted sites that cannot ever access it. (Incidentally, the IETF's Geopriv work has a built-in whitelisting capability.) That way, users could avoid the 24-hour permission renewal described above and they would not be badgered into consenting by accident.

This kind of granularity would also help with permission revocation. Right now, to revoke even a single site's permission, the only choice is to revoke all sites' permissions. Even accomplishing that is a counterintuitive process: under the general settings, using the tab marked "Reset" (a somewhat scary name), the user must select "Reset Location Warnings." Granted, the 24-hour permission relapse means that, today, there probably will not be many sites to revoke permissions from. But if the permission model ever changes, the revocation model needs to change as well.

Given the privacy interests at stake and the relative lack of protection in the law, we would expect location controls to be better than other kinds of technological controls on the Web, to offer users more choices about what happens to their data and to be especially transparent about when location data is being passed around. It does not appear that every one of our expectations will be met here at the dawn of the location-enabled Web. But as location comes to pervade the Web experience - which it will, given the simple interface offered by the browser vendors and myriad uses of location information - we will be taking a closer look at how current user controls work, how they could be improved, and how standards, policy, and law can contribute to protecting location privacy on the Web.

IETF Geopriv Working Group [6] (February 2009)

Draft WC3 Standard [7] (June 2009)

---

3) Location-Aware Firefox

Firefox, the second-most popular Web browser in the US after Microsoft's Internet Explorer, has also recently become location-enabled. As with Safari on iPhone, this means is that Web sites can now ask Firefox for your location, and the browser can now deliver it. Initially, Google has signed on as the default "location provider" for Firefox. When a Firefox user pulls up a Web site that wants to use his or her location, Firefox will gather some information about nearby WiFi access points and send that information to Google. Because Google maintains a database that maps WiFi access points to actual physical locations, it can use this information to calculate the user's location. That location gets sent back to the Firefox browser, and the browser forwards it on to the Web site that originally requested it. The accuracy of the location depends on a number of factors, but can be within a handful of meters in densely populated areas.

Firefox and Google have taken a couple of excellent steps to protect the privacy of Firefox users throughout this process. The location information gets transmitted over an encrypted connection so it cannot be sniffed en route between the browser and Google or vice versa. Firefox does not provide Google with any information about the site that made the location request, so Google does not learn anything extra about the user's browsing habits. Google also de-identifies the information it receives from Firefox two weeks after it is collected.

This seems like a solid set of standards that all location-enabled browsers and location providers should be able to meet. While it is nice to see Google and Firefox take these steps, we are hopeful that Firefox will be able to expand its pool of location providers, and that new location providers will be able to meet these same standards. There are actually a diversity of ways in which Web users can or will soon be able to obtain their own locations, and as new location providers crop up, users should have the ability to choose their preferred provider.

On the user experience side, the story is somewhat mixed. While Firefox, like Safari on iPhone, will prompt users for permission before passing location on to a Web site, there is no easy way to see a list of sites that have obtained location. If the user loses trust in a particular site, he or she must go back to the site itself to revoke its permission, which is probably precisely what the user will not want to do. And the mechanism for disabling location-awareness altogether is somewhat complex. We expect to see some more intuitive user controls for these kinds of features as more and more Web sites become location-enabled.

Geolocation in Firefox [8] (June 2009)

Links:
[1] http://cdt.org/publications/policyposts/2009/12#1
[2] http://cdt.org/publications/policyposts/2009/12#2
[3] http://cdt.org/publications/policyposts/2009/12#3
[4] http://www.cdt.org/privacy/20090423_dpi_testimony.pdf
[5] http://www.cdt.org/publications/digital-search-and-seizure.pdf
[6] http://www.ietf.org/html.charters/geopriv-charter.html
[7] http://dev.w3.org/geo/api/spec-source.html
[8] http://www.mozilla.com/en-US/firefox/geolocation/