From: Electronic Frontier Foundation Jan. 5, 2009
To: Federal Trade Commission
Re: Privacy Roundtables – Comment, Project No. P095416

Contents

The Modern Web Offers Users Little or No Privacy....1
Trackability is a natural by-product of sophisticated hypertext systems....1
There are more than ten mechanisms available for tracking web users....1
The economics of the Web encourage tracking....2
Complexity and lack of transparency preclude a functioning market for privacy....2
Existing Privacy Enhancing Technologies are Inadequate....4
Most PETs only address a small subset of the 10+ tracking mechanisms....4
Of the few PETs that afford significant protection, most are awkward and difficult to use....4
Nobody has adequate economic incentives to develop high-quality PETs....5
Browser manufacturers have no incentive to distribute the most effective PETs....5
Finding the Right Role for Privacy Regulation....6
Would strong privacy regulation pose a threat to innovation?....6
What would good privacy regulation look like?....6

The Modern Web Offers Users Little or No Privacy

Trackability is a natural by-product of sophisticated hypertext systems

In a fundamental sense, the features required for sophisticated hypertext systems also make for highly effective surveillance systems. For instance:

- The ability to request a hypertext document from a remote computer requires that you send a reply IP address to that remote computer, normally allowing the server to see which computers are fetching which documents.
- The ability for a page on server A to embed text, images, video or other content from server B also enables "web bugs" that exist solely as a way for server A to cause visitors' browsers to report each visit to B.
- The existence of cookies, which are necessary for web sites to offer stateful user interfaces, also provides a means for sites (and third-party profilers) to track their users across many visits and across the IP addresses at their home, work, cafes, and travel destinations.
- Client-side scripting systems, such as JavaScript, Java, Flash and Silverlight, are necessary for the more dynamic behavior that some modern websites display, but these systems also provide a way for pages to probe many aspects of the user's behavior, and to "phone home" reports of these probes.

There are more than ten mechanisms available for tracking web users

The culmination of the developments discussed above is that today, there are more than ten features commonly or universally built into web browsers that are either (a) capable of being used to track that browser, or (b) capable of being combined with one or two other features in order to produce a tracking system. These mechanisms are:

1. IP addresses;
2. cookies;

3. accounts on web servers[1];
4. User Agent strings, plugin versions and sundry browser characteristics;[2]
5. embedding of third party objects;
6. scripts that automatically transmit data to 2nd or 3rd parties;

and at least five kinds of "supercookies" created by:

7. Adobe Flash;
8. HTML 5 DOM Storage;[3]
9. Microsoft Silverlight;
10. Google Gears; and
11. Internet Explorer userData.

The economics of the Web encourage tracking

For websites and other online service providers, there are financially significant upsides to tracking their users, and few downsides. Only a small proportion of their users will read a privacy policy, and only a much smaller proportion will understand what it actually means.[4] On the upside, websites can understand their userbases and target advertisers at them more profitably when equipped with complete histories of each visitor's activities. On most occasions when loss of privacy causes harm, service providers are safe from liability, even when their design decisions and policies made the loss of privacy possible.

However strong the incentives may be for websites to track their users, the incentives are even stronger for advertising and tracking networks. These companies have no relationship whatsoever with the people whom they surveil and collect records of, and their privacy policies thereby escape the attention of the people affected by them. These firms compete with one another to offer advertisers the most detailed information about the character and history of the people that advertisements were displayed to, and which of those people were interested in any given product. Thus, the marketplace for web advertising will produce increasingly intrusive architectures for user profiling regardless of whether such intrusions serve a sufficiently useful social purpose to outweigh their harmful side effects on privacy.

Complexity and lack of transparency preclude a functioning market for privacy

Many participants in debates about online privacy agree that we should have marketplaces in which users can make informed choices between services, based not only on the cost of each service but

---

1 Accounts on web servers are often, but certainly not always, tracked using cookies. In general, account functionality is implemented using some mixture of HTTP POST requests, HTTP authentication, cookies, explicit session IDs in URLs, or even hidden session IDs passed using POST requests.

2 Other browser characteristics which can be combined to distinguish between one machine and another include timezone, screen resolution and color depth,

3 To investigate whether websites are already using non-Flash supercookies in the wild, EFF performed an informal survey of the use of DOM storage. After visiting about 30 major US websites, one 2nd party and one 3rd party had set persistent DOM storage supercookies on our test system. The phenomenon deserves further study along the lines of Soltani *et al*.'s work on Flash LSO objects, *Flash Cookies and Privacy* http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1446862 .

4 Research by Cranor and McDonald found that users believe that privacy policies mean that they have legal protection. They observed: "if people hold the mental model that any company with a privacy policy is bound by law not to release data to third parties, and if that is the only threat that worries them, why would people bother to read the policy?" *An Empirical Study of How People Perceive Online Behavioral Advertising*, CMU-CyLab-09-015 (Nov. 10, 2009), available at <http://www.cylab.cmu.edu/files/pdfs/tech_reports/CMUCyLab09015.pdf> (p. 4 & n. 16)

on the impact it will have on their privacy. In other words, each account a consumer could sign up for could have two prices: a price in dollars and a price in privacy. Unfortunately, the marketplace does not offer consumers any practical way to choose between services based on their prices in privacy.

The structural separation discussed above between websites (which may have a business relationship at least with those users who choose to create accounts on them) and tracking and advertising firms (which have no relationship with the people they track) is an example of institutional complexity that reduces the practicality for consumers of choosing between services based on their privacy practices.

Similarly, technical complexity can reduce the practicality of choosing between services based on privacy. When Google, Microsoft and Yahoo! announced that they decrease their periods of retention for partial IP address logs and cookie data, most consumers had no way to answer the obvious question: will this search engine still have an identifiable record of my search after a period of *x* months?

Even when online firms do their best to provide concise and readable descriptions of their privacy practices, one is typically left with one to two pages of subtle and nuanced English text to read. When multiplied by the typical number of accounts that a typical American Internet user maintains, the result is a volume of information that only professional privacy experts have much hope of understanding, let alone comparing to the other available service in the market.

This problem naturally leads to the subject of privacy enhancing technologies, since that field holds out the hope that individuals could enter *their* privacy policies in one place, and hold the rest of the Internet to that policy.

Unfortunately, as we will see, privacy enhancing technologies are not yet ready to deliver this vision.

Existing Privacy Enhancing Technologies are Inadequate

Most PETs only address a small subset of the 10+ tracking mechanisms

If a user searches for software or browser plugins to improve their web privacy, the most likely outcome is that they will find a tool that promises a lot but does not systematically address all of the mechanisms by which browsers are routinely tracked.

Even what we believe to be the state-of-the-art PETs do not address all of the problems at once, and probably have to be used in combination with one another to be effective:

	Proxy / VPN	Tor & TorButton	RequestPolicy	NoScript	TACO	AdBlock Plus
IP address	Yes	Yes	No	No	Rarely	No
Cookies	No	Often	No	Rarely	Rarely	No
Accounts	No	No	No	No	No	No
User Agent + plugins	Rarely	Yes	No	No	Rarely	No
3rd party content*	No	No	Yes*	Only scripts	Sometimes *	Often*
Scripts	No	Yes	No	Yes	Rarely	No
Flash Cookies	No	Yes	No	Yes	Rarely	No
DOM storage	No	Yes	No	Yes	Rarely	No
Silverlight	No	Yes	No	Yes	Rarely	No
Gears	No	Yes	No	Yes	Rarely	No
IE userData	No	N/A	N/A	N/A	N/A	N/A
Convenient to use?	Yes	No	No	No	Yes	Yes

* 3rd party content can encompass any of the other tracking mechanisms if that mechanism is employed via a 3rd party server, so a PET that prevents tracking by 3rd parties may partially address any or all of the other tracking methods.

Of the few PETs that afford significant protection, most are awkward and difficult to use

As explained above, most of the tracking mechanisms built into web browsers appeared as a result of efforts to expand the functionality of hypertext systems: rather than sitting down to design tracking systems, engineers sat down and designed features that happened to double as tracking mechanisms.

This means that the easiest way for a PET to adequately protect a user from being tracked via some feature of hypertext is to simply disable the feature in question: turn off cookies, do not load third party content, do not execute scripts, etc. Unfortunately, taking this road sometimes leaves users surfing the web as though it were 1990 all over again. Some better-designed sites continue to work, but many “flashier” websites simply cease to function.

A common response by PETs to *that* problem is to allow users to re-enable sophisticated hypertext features on a site-by-site basis: reenable session cookies, 3rd party content, scripts, etc. Unfortunately, the result is an interface that is hard to navigate and non-intuitive: it's hard to tell when a website isn't functioning because it needs a cookie, because it needs javascript, or when it simply isn't compatible with your browser version.

Sometimes PETs do a better job at automatically guessing which uses of a hypertext feature are in the interest of a user, and which are simply intended to track them. AdBlock Plus deserves particular commendation in that regard. Unfortunately, no currently existing PETs seem to achieve that for all or even most of the 10+ tracking mechanisms!

Nobody has adequate economic incentives to develop high-quality PETs

Of all the genuinely useful PETs listed in the table above, only one, Tor, is developed by an organization that has anything resembling systematic funding for its work. The others are developed by volunteers, or by academic researchers who are not paid for producing working, reliable software.

Tor has been able to secure *sui generis* grants to support its work from a range of eclectic sources including military research institutions, private organizations interested in supporting freedom of speech,[5] and US government agencies interested in offering censorship-resistant communications tools to foreign nationals. But the existence of *sui generis* grants for one PET does not indicate that there are systematic economic incentives for the development of PETs – if anything, it indicates the opposite.

Key intermediaries have no incentive to distribute the most effective PETs

Given how limited their funding is, PET developers cannot be expected to have their own extensive marketing and distribution campaigns. Thus, the public is unlikely to use any given PET unless it is bundled by default with common web browsers. Browser manufacturers thus serve as key intermediaries necessary for the mainstream adoption of PETs. But it is unclear whether browser makers have much incentive to embed or promote PETs.[6] In addition, most common browsers are aligned with entities with ambiguous incentives as to user privacy. Internet Explorer is a Microsoft product; Google produces Chrome and provides over 90% of the funding for Mozilla, developer of Firefox.[7] It is unrealistic to expect browsers produced or funded by such actors to aggressively incorporate strong privacy protections if advertisers remain hostile to these protections.

Web browsers do engage in limited PET development and competition over privacy features. For instance, all of the major browsers have recently added some kind of "private mode" that allows users to open tabs that are theoretically "separate" from other uses of the browser.

But browser private modes are mostly directed and engineered to protect users against tracking by other people with access to their computer. For instance, browser private modes are effective at keeping visited websites out of a browser's local history logs. By contrast, these tools are less effective for preventing tracking by remote websites. For instance, when researchers tested them, none of the browser private modes were effective at preventing tracking by all of the kinds of

5 In the interest of full disclosure, we should note that EFF provided interim funding to the Tor project while it was establishing itself as an independent nonprofit. Other nonprofit foundations have subsequently offered grants and assistance for Tor development.

6 For instance, hundreds of users since 2001 have asked the Mozilla developers to add a complete ad-blocking system to Firefox https://bugzilla.mozilla.org/show_bug.cgi?id=94035. And since the development of AdBlock Plus, such a system even exists and Mozilla merely needs to ship it with their browser.

7 See http://www.mozilla.org/foundation/documents/mf-2008-audited-financial-statement.pdf ("Approximately 91% and 94% of Mozilla revenue for 2008 and 2007, respectively, was derived from this [Google] contract.")

supercookies.[8] None of the browser private modes do anything to inhibit tracking via the combination of IP address, User Agent, and plugin version information.

Browser manufacturers could easily decide to ship PETs that were more effective than their existing privacy modes. For instance, they could add a setting so that each browser tab or window was a separate private browser session (preventing, for instance, a gmail login in one tab from using the same cookies as a google search in another tab). Or browser manufacturers could include 3rd party PETs that are both effective and easy to use, such as AdBlock Plus or TACO. But they have chosen not to do so.

In other fields of technical innovation besides the Web, the structural problem of key intermediaries who are not particularly committed to privacy is even greater. In the

Finding the Right Role for Privacy Regulation

Would strong privacy regulation pose a threat to innovation?

The roundtable questions ask whether consumer privacy can be protected without stifling innovation. We should not presuppose that there is a dichotomy between the social good of privacy and the social good of innovation. In fact, the relationship between the two is more complicated.

Some innovation has clear negative impacts on privacy, while other innovation may be privacy neutral or even privacy enhancing. Firms need stronger incentives to do privacy-neutral and privacy-enhancing innovation and fewer incentives to do innovation that is harmful to privacy.

A properly functioning and transparent market for digital services would provide those kinds of incentives for privacy-neutral and privacy-enhancing design. But, for the reasons discussed above, we don't have one.

What would good privacy regulation look like?

Discussions about regulatory solutions to online privacy problems usually revolve around the codification of some variant or subset of the Fair Information Practices.

There is no question that making the Fair Information Practices enforceable would be a substantial step forward. There are, however, several areas of concern. One important concern is the sheer volume of reading required to track even comparatively brief and clear privacy policies for all of the services that a typical Internet user interacts with. Another is the need to find a good way to implement the principle of access efficiently and securely for online services.

These problems may be solvable with regulatory innovation. But in our view, ideal privacy regulation would go beyond the Fair Information Practices, and strike at the underlying lack of incentives for the development, deployment and support of privacy enhancing technologies.

We look forward to discussing what such regulatory innovation would look like at the FTC's privacy roundtables.

8 Katherine McKinley, *Cleaning up after cookies,* iSec Partners White Paper, https://www.isecpartners.com/files/iSEC_Cleaning_Up_After_Cookies.pdf Section 2.2