

**Original Article**

# **REVEALING THE DEPTHS OF THE WEB: AN EXTENSIVE INVESTIGATION INTO THIRD-PARTY HTTP REQUESTS ACROSS ONE MILLION WEBSITES**

*Elena Martinez*

Faculty of Information Science, University  
of Barcelona, Spain

**Abstract:** In the era of pervasive information technologies, the landscape of modern society has been fundamentally reshaped by their remarkable benefits, yet it has also witnessed the emergence of extensive surveillance practices. Alan Westin, half a century ago, foresaw the potential "intolerable" nature of information technology's surveillance applications. Today, it is widely acknowledged that every new digital device, application, and social platform carries concealed privacy risks. In contrast to the cash-based transactions of the previous century, personal behavioral data has evolved into a new form of currency in the current era.

Understanding individuals' online behaviors, whether it involves website browsing habits or product preferences, has become instrumental in precisely targeting products and services to individuals with specific behavioral profiles. This transformation raises profound questions about the balance between technological advancements and privacy rights in the digital age.

**Keywords:** information technology, surveillance, privacy risks, personal data, behavioral profiling.

## **Introduction**

As revealed over the past several decades, the beneficial aspects of information technologies have proven truly remarkable and fundamentally transformative, yet they also have ushered in an age of nearly limitless surveillance. It is this use of information technology for surveillance purposes that Alan Westin described as potentially "intolerable" nearly 50 years ago (Westin, 1967,). Indeed, every new device, app, and social network is now assumed to come with hidden privacy risks. Whereas the economy of the past century relied on cash for most transactions, personal behavioral data has become a currency of its own in the current century. Learning what an individual does with her or his time, be it browse specific websites or purchase particular products, has become a means to more efficiently target products and services to those who fit certain behavioral profiles.

Efforts to collect personal behavioral data have extended to governments as well. News reports based on internal documents of the United States National Security Agency (NSA) leaked by former contractor Edward Snowden have revealed that the U.S. government has constructed a surveillance system of such complexity and scope that it fulfills the dire warnings made by Senator Frank Church, who suggested in 1975 that the NSA's technological capabilities carried with them "the capacity . . . to make tyranny total in America," and that a fully realized technological surveillance state constituted an "abyss from which there is no return" (Bamford, 2011, p. 2). Looking back on Westin's and Church's worries in the context of the present day, we understand that their warnings, however extreme, had significant merit. It is now clear that information technologies carry with them externalities that are socially intolerable, and that the realities of state surveillance

## **Original Article**

often exceed the worst predictions of yesteryear. However, despite the warnings of Westin and Church proving prescient, many opportunities exist to shift the balance of power in favor of privacy and user choice.

Westin and Church both dealt with many hypothetical scenarios and little knowledge of future technologies, but today the general principles of computing, data storage, and networking are well known. Those with technical knowledge of computing systems may use that knowledge to analyze the surveillance systems that affect the lives of billions of Internet users. As with discussions of other social ills such as poverty and preventable disease, debates on issues of surveillance are served when backed by sound quantitative measures. Thus, to escape Church's "abyss," public policy advocates must identify and enumerate specific forms of surveillance, the scope of this surveillance, and the commercial and government practitioners of it.

The purpose of this article is to provide a quantitative analysis of privacy-compromising mechanisms on the top 1 million websites as determined by Alexa. It is demonstrated that nearly 9 in 10 websites leak user data to parties of which the user is likely unaware; more than 6 in 10 websites spawn third-party cookies; and more than 8 in 10 websites load Javascript code. Sites that leak user data contact an average of nine external domains. Most importantly, by tracing the flows of personal browsing histories on the Web, it is possible to discover the corporations that profit from tracking users. Although many companies track users online, the overall landscape is highly consolidated, with the top corporation, Google, tracking users on nearly 8 of 10 sites in the Alexa top 1 million. Finally, by consulting internal NSA documents leaked by Edward Snowden, it has been determined that roughly one in five websites are potentially vulnerable to known NSA spying techniques at the time of analysis. It should be noted that this is not the first quantitative study of privacy online. However, the vast majority of quantitative research in the area of Web tracking has appeared in the computer science literature, with scholars such as Krishnamurty and Wills conducting foundational studies as far back as 2006. These studies provide essential insights into the nature and scope of online tracking, but they require a high level of technical expertise to understand. This has had the unfortunate side effect of reinforcing barriers between disciplines and reducing the potential for cross-disciplinary work. This is despite the fact the communication discipline has traditionally played a prominent role in areas of telecommunications and media policy. For these reasons, this article aims to bridge a disciplinary gap.

This article begins with a technical primer on how Web tracking works, explores previous research in the area, and then explicates several research questions. This is followed by a section that details a new open-source software platform named webXray that has been developed to conduct this research. Finally, findings are presented along with suggestions for how policy makers may shift the balance away from unchecked surveillance toward respect for user preferences.

### **Background: Third-Party HTTP Requests and User Tracking**

To fully appreciate how companies track users on the Web, a technical aside is in order. When a user's Web browser makes a request to load a web page, it receives an HTML file from the server associated with a given domain (such as "http://example.com," where "example.com" is the domain). Any request to download content, such as a web page, uses the hypertext transfer protocol (HTTP) to connect the user to the server hosting the desired content (Fielding et al., 1999). Only two parties are involved in the transaction: the user and the company that owns server. Due to the fact that only one server is being connected to, and the address of that server is visible in the browser address bar, these are often called *first-party* requests.

The beauty of HTML, and the Web itself, is that HTML pages can include additional pieces of content, called *elements*, from a limitless array of additional parties. Such elements may include pictures, videos, and Javascript code. However, when new content is included from outside domains, the number of parties involved increases from two to three: the user, the site she or he is visiting, and the entity providing the additional content. For this

## Original Article

reason, an image or other piece of content from an external server may be called a *third-party element* (sometimes abbreviated to 3PE), and the process of downloading such an element is the result of a *third-party request*. Every time an HTTP request is made, information about the user is transmitted to the server hosting the content. These data include the IP address of the computer making the request, the date and time the request was made, and the type of computer and Web browser employed by the user, which is known as the *user-agent* field. In addition, the address of the page that initiated the request, known as the *referrer* [sic], is included. This information is found in the raw data of the HTTP request itself, an example of which is shown here:

IP: 8.67.53.09

DATE: [10/May/2014:19:54:25 +0000]

REQUEST: "GET /tracking\_pixel.png HTTP/1.1"

REFERER: "http://example.com/private\_matters.html"

USER-AGENT: "Macintosh; Intel Mac OS X 10\_8\_5...AppleWebKit/537.71"

From the above information, the server receiving the request for the file called tracking\_pixel.png can determine that it was made from a user with the IP address 8.67.53.09, using the Safari browser on a Macintosh computer, who is currently viewing a web page with the address http://example.com/private\_matters.html. If the server has many such records, patterns of behavior may be attributed to the same combination of IP and user-agent information. This is the most basic form of tracking and is common to all HTTP requests made on the Web.

Before companies engaged in Web tracking may harvest HTTP request data, they must first get their elements included in the source code of a page. In some cases, advertisers pay to have their elements on a page, and this is the dominant business model for online publishers (Turow, 2011). However, most third-party HTTP requests are made to corporations such as Google, Facebook, and Twitter, which provide free services such as traffic analytics and social media tools. Although webmasters do not pay for such tools, the hidden cost of privacy loss is incurred by unwitting users.

Despite the fact that this situation is already rather murky from the user's perspective, it actually gets worse. Some third-party elements are also able to store information on a user's computer using small files called *cookies* (Kristol & Montulli, 1997). Cookies may be used as a unique identifier that is affixed to the browser of a given user. This is analogous to the tracking bracelets ornithologists place on migratory birds; even if users visit unrelated sites, tracking cookies have rendered them tagged and traceable. A user may delete or block cookies, potentially restoring some privacy. However, other third-party requests may download sophisticated Javascript code that can execute commands on users' computers, facilitating advanced forms of tracking known as *browser fingerprinting* (see Eckersley, 2010, for an early study on the topic). By detecting the unique characteristics of a given computer, browser fingerprinting is analogous to its off-line form: The specific list of fonts installed on a computer may function like a loop or a whorl on a finger—given combinations are often unique to an individual. Also like real fingerprints, it is often difficult to hide or change a browser's fingerprint. The attention of corporations is rapidly turning to fingerprinting techniques for this reason, and cookies are quickly falling out of favor.

## Previous Research

Previous research in Web privacy has generally focused on three main areas: attitudes of users toward the idea of being tracked, the mechanisms which facilitate tracking, and measurement of how much tracking is happening. The majority of research regarding attitudes has taken place within the fields of communication and law, whereas mechanisms and measurement research has primarily been the domain of computer science.

### Attitudes

## **Original Article**

Opinion research regarding online privacy has been ongoing since the 1990s. Yet, as much as the Web has changed since that decade, attitudes toward privacy have remained fairly constant. In 1999, “respondents registered a high level of concern about privacy in general and on the Internet” (Ackerman, Cranor, & Reagle, 1999, p. 2). Likewise, a 2003 study found that “a clear majority of Americans express worry about their personal information on the web” (Turow, 2003, p. 16). A 2009 survey found that “69% of American adults feel there should be a law that gives people the right to know everything that a website knows about them” (Turow, King, Hoofnagle, Bleakley, & Hennessy, 2009, p. 3). Similar work in 2012 showed that 60% of respondents would like the proposed Do Not Track (DNT) standard to prevent websites from collecting personal information (DNT is further explained in the Discussion section) (Hoofnagle, Urban, & Li, 2012). A 2014 Pew Research Center study titled “Public Perceptions of Privacy and Security in the Post-Snowden Era” discovered that 70% of respondents felt records of the websites they had visited constituted “very” or “somewhat” sensitive information (Madden, Rainie, Zickuhr, Duggan, & Smith, 2014). One reason marketers may argue that online data collection is acceptable is that consumers are making a trade-off between privacy and discounts; yet a 2015 study found that 91% of respondents *disagreed* with the statement “If companies give me a discount, it is a fair exchange for them to collect information about me without my knowing it” (Turow, Hennessy, & Draper, 2015, p. 3). The general trend established by these surveys is not just that most users value privacy but that the current state of privacy online represents an area of significant anxiety. Indeed, research into the mechanisms of tracking highlight that there is serious cause for concern.

### ***Mechanisms***

Computer security research has explored many novel methods of compromising user privacy. These have included detecting whether a user had visited a page based on the color of a given hyperlink (Jackson, Bortz, Boneh, & Mitchell, 2006), analyzing Flash cookies that survive user deletion (Ayenson, Wambach, D., Soltani, A., Good, N., & Hoofnagle, 2011), the increasingly popular fingerprinting techniques discussed above (Acar et al., 2013; Eckersley, 2010; Nikiforakis et al., 2013), as well as other uses of Javascript (Jang, Jhala, Lerner, & Shacham, 2010). The common thread among such research is a focus on ways in which a browser may be used to extract more information about the user than is typically provided by HTTP request headers. This focus on what is called the client side (i.e., the user’s computer rather than the server) is born of the fact that researchers have long assumed “that sites are not able to reliably track users using just their IP address and user-agent string” (Jackson et al., 2006, p. 740). In other words, the basic information detailed above from an HTTP request is insufficient to track users with the same level of accuracy as other methods such as cookies or fingerprinting. However, this assumption has been proven wrong.

Even without direct browser manipulation, Yen, Xie, Yu, Yu, and Abadi (2012) have demonstrated that solely by analyzing HTTP headers on a multimillion-point data set, they were able to reach a user identification rate of 80%, which is “similar to that obtained with cookies” (p. 2). This finding is significant because it demonstrates that third-party HTTP requests by *themselves* hold potential for compromising user privacy similar to the more commonly studied client-side techniques. A single HTTP request is unlikely to result in user identification, but thousands of such requests may be correlated to a given user. Such techniques happen on the servers of corporations and are largely hidden from analysts of client-side code.

### ***Measurement***

The third general area of research focus has been in measurement. In this area, researchers are interested in selecting a population of websites and essentially conducting a census of tracking mechanisms, third-party requests, and data flows. Much of the pioneering work in this area has been conducted by Krishnamurthy and Wills. In 2006, they advanced the concept of the “privacy footprint,” which they describe as a means to analyze the “diffusion of information about a user’s actions by measuring the number of associations between visible



## **Original Article**

nodes via one or more common hidden nodes” (Krishnamurthy & Wills, 2006, p. 70). What Krishnamurthy and Wills describe as “hidden nodes” may be more widely included under the umbrella of third-party requests. This work was followed up a few years later with a longitudinal study (Krishnamurthy & Wills, 2009). The data presented in this article update and advance Krishnamurthy and Wills findings with a larger sample.

Although Krishnamurthy and Wills have done an excellent job in the area of measurement, they are far from the only researchers exploring the topic. Castelluccia, Grumbach, and Olejnik (2013) have recently analyzed “the flows of personal data at a global level” (p. 1) and found intriguing differences in the nature, scope, and ownership of tracking mechanisms around the world. Other researchers have developed desktop browser add-ons (Mayer & Mitchell, 2012; Roesner, Kohno, & Wetherall, 2012), which they have used to detect the presence of various tracking mechanisms. Some of the most recent measurement literature focuses on fingerprinting specifically (Acar et al., 2013). Additional research has revealed extensive tracking on websites containing health information (Libert, 2015). A common theme among all measurement research is that the amount of tracking on the Web is increasing and shows no signs of abating.

### **Research Questions**

This article seeks to update and advance knowledge in three main areas: the general size and nature of the hidden Web, the corporate ownership patterns that affect user privacy, and the degree to which state surveillance is facilitated by commercial tracking mechanisms. In regard to general questions on the size and scope of the Web tracking ecosystem, several issues are explored in depth, including the number of pages that leak user data, how many entities users are exposed to in tandem, and how this varies among different countries. Furthermore, third-party elements are investigated to determine whether they are visible to the user or represent potential attempts at covert tracking. Second, and of central interest to this study, the corporations that are responsible for tracking users online are discovered. This focus sheds light on the larger social and political-economic considerations involved in Web tracking. Finally, internal documents released by former NSA contractor Edward Snowden have revealed that the spy agency has leveraged commercial tracking tools to monitor users. Thus, this research project also seeks to detect commercial tracking mechanisms known to be co-opted by the NSA.

### **Method**

Following the findings of Yen et al. (2012), it is clear that third-party HTTP requests provide an excellent unit of analysis for examining the extent of tracking on the Web. Such requests are potentially as revealing and invasive as cookies, and they are a basic prerequisite of every single family of tracking techniques. Regardless of what cookies are set or what fingerprints are taken from a browser, a connection between the user and the server of the tracker needs to be established via HTTP. By examining HTTP requests in bulk, a picture of data flows emerges. On a sufficiently large data set, those requests that have a benign function fade into the long tail of data, and those that happen with great regularity are easily spotted. Therefore, this research faces two basic problems: determining which sites form a sufficiently large population and how to analyze third-party HTTP requests on each site.

The best population of sites was determined to be the top 1 million sites list published by Alexa. Alexa is a subsidiary of Amazon that provides website traffic metrics and rankings derived from a “panel of toolbar users which is a sample of all internet users.” The degree to which Alexa’s data accurately represent popularity on the Web is debatable, because Alexa Toolbar users are by no means a random sample of all Internet users. However, it has become common practice for researchers to use the Alexa list in the absence of other, better lists (Acar et al., 2013; Castelluccia et al., 2013; Jang et al., 2010; Krishnamurthy & Wills, 2006, 2009; Mayer & Mitchell, 2012; Nikiforakis et al., 2013; Roesner et al., 2012). For this study, the full Alexa list was downloaded in May 2014.

To detect tracking on the sites selected, the webXray software (version 1.0) was used. webXray is written primarily in Python and works in several steps. First, a list of website addresses is given to the program. This list is processed to ensure all the addresses are properly formatted, are not links to common binary files (such .pdf or

## Original Article

.xls), and are not duplicates. Next, each site is relayed via the Python *subprocess* module to a command-line instantiation of the headless Web browser PhantomJS. In this context, *headless* refers to the fact that the browser runs on a command line and does not require a graphical user interface; this makes PhantomJS ideal for deployment on a cloud-based virtual machine. PhantomJS takes as arguments a Web address and a Javascript program, which are responsible for loading the website, processing the page title and metadata, collecting cookies, and detecting both HTTP requests and HTTP received events. To successfully follow multiple redirects and allow elements to download, PhantomJS is given 30 seconds to complete loading the page. The results are passed back to the Python script as JSON data and processing resumes. Data about the page, such as title and meta description, are stored in a database, and then HTTP requests are examined.

webXray uses the Mozilla Public Suffix List (<http://publicsuffix.org/>) to determine the domain of each HTTP request (subdomains are ignored). Requests that have the same domain as the page address are ignored (e.g., “<http://example.com>” and “<http://images.example.com/header.png>”), whereas requests with different domains are entered into a database of third-party elements (e.g., “<http://example.com>” and “[http://www.google-analytics.com/\\_utm.gif](http://www.google-analytics.com/_utm.gif)”). This process is repeated for storing third-party cookies. Request strings are additionally parsed to remove argument data from the domain and element. For example, the following string:

[http://sub.example.com/tracking\\_pixel.png?id=8675309](http://sub.example.com/tracking_pixel.png?id=8675309) is broken up as having the domain “example.com,” the element “tracking\_pixel.png,” the arguments “?id=8675309,” and the file extension “png.” File extensions of elements (e.g., “png,” “js,” “css”) are used to catalogue elements by type (e.g., “image,” “Javascript,” “Cascading Style Sheet”).

Finally, after a scan is completed, the top 100 domains are identified and ownership information is manually determined. The process primarily relies on using the *whois* utility to look for domain registration records. In cases where whois records are unhelpful, the domain or element is loaded into a browser to follow redirects to a parent site. If this fails, Web searches are used to find ownership. When smaller companies are found, the website Crunchbase is consulted to determine whether they are subsidiaries or acquisitions of larger companies. Great care in this process is taken to ensure ownership findings are accurate and the “About” pages of all companies are consulted as part of the process. One additional barrier is the fact that many companies have several domains. For example, Google has by far the most domains, with several hundred unique entries listed on Wikipedia ([https://en.wikipedia.org/wiki/List\\_of\\_Google\\_domains](https://en.wikipedia.org/wiki/List_of_Google_domains)); the names of these domains range from the transparent (google.com, google.fr) to the opaque (1e100.net, ggph.com). Adobe was found to be using many domains as well—again, some obvious (adobe.com) and some not (2o7.net). This pattern of household brand names using unfamiliar domains held fast for several companies, including Twitter, Facebook, and Amazon. It should be noted that there is nothing particularly nefarious about companies using unfamiliar domains to host content, only that it presents a barrier during the analysis process. In total, more than 500 distinct domains were traced back to 140 different companies. To reduce falsepositives, domains that are requested by only a single site are ignored in the final analysis. For example, if the site <http://example.com> was the *only* site to initiate a request to the domain example-images.com, it is ignored. This approach leaves only domains that can conceivably link visitors between two or more sites in the analysis pool.

## Findings

### General Trends

The home pages of the most popular 990,022 websites as measured by Alexa were analyzed in May 2014. Ninety-six percent of sites (950,489) were successfully analyzed, and 4% failed to load. This failure may be attributed to the top sites’ lists containing sites that were not available at the time of analysis (other webXray tests yielded success rates of over 99%, indicating the software is not at fault). The pages analyzed yielded a set of 7,564,492 unique cookies and 21,214,652 unique element requests and encompassed 1,056,533 distinct domains. The

## **Original Article**

number of pages that initiated requests to third parties was 832,349, representing 88% of the total. Further analysis determined that sites that *do* make requests to third parties contact 9.47 distinct domains on average—indicating that user data are not only being frequently shared but often being shared with many parties simultaneously. 62.9% of sites spawn thirdparty cookies, which may be used to track users using traditional methods, and 82.99% of sites include third-party Javascript, which may represent the increasing trend toward fingerprinting and other advanced techniques. It should be noted that because this study used a specified population rather than a statistical sample, the confidence for these numbers is 100%, but these results *only* represent the Alexa top 1 million, *not* the entire Web.

In addition to the top-level trends, subanalyses were conducted on the 10 most common toplevel domains (TLDs) as well as U.S. educational (edu) and government (gov) sites. The top TLDs include the well-known com, net, and org as well as several country-code top-level domains (ccTLDs) such as Russia (ru), Germany (de), United Kingdom (uk), Brazil (br), Japan (jp), Poland (pl), and India (in). Sites ending in a ccTLD have been registered within a given country; however, popular sites within a given country often include com and org sites, so this is not a perfect proxy for what a user from a given country is exposed to (i.e., do not assume Russians only visit ru sites). When comparing data among TLDs, there was not much range in the percentage of sites that had a third-party element: Russian domains had the most with 94.33%, and Indian domains had the least with 81.79%. However, in regard to the average number of domains contacted, there was a large spread, with Brazilian sites contacting 11 on average and U.S. government sites contacting 3.61. There was also a wide disparity in the percentage of sites spawning third-party cookies: Russia was again in first place with 89.42% of sites spawning cookies, whereas only half as many U.S. government sites (39.63%) did so. Finally, there was not much difference in the percentage of sites with third-party Javascript: At the high end were Russian sites with 88.72%, and at the low end were Indian sites with 77.14%. These findings are presented in greater detail in Table 1.

### ***Types of Requested Elements***

To determine current trends in tracking mechanisms, the top 100 requested elements were categorized according to six major types based on file extensions. Javascript is by far the most prevalent, being requested in 36% of cases. Javascript is a popular development language, and many uses of Javascript are wholly benign, yet still generate HTTP request records and potentially may be used for advanced fingerprinting techniques. The analysis of Javascript code is outside the scope of this article, and this prevalence is noted to highlight trends, not to suggest all Javascript is designed for fingerprinting. Nevertheless, Javascript elements may pose a greater potential privacy risk than inert elements such as images. Pages with the ability to serve dynamic content, such as PHP and CGI scripts, make up 9% of the total. Images, such as tracking pixels, have a history stretching back to the early days of the Web and are still popular, comprising 7% of the total. Fonts and Cascading Styles Sheets together make up 5% of requests. Two requests for structured data (JSON) were found in the top 100. Thirty-eight percent of requested elements were not categorizable because they lacked recognizable file name extensions.

### ***Images***

Of the seven identified element types, six are in no way visible to the user and provide no clue that she or he is being tracked. However, images could presumably offer a means for a user to determine that additional parties are present. Yet, as detailed in Table 2, only 4 of the top 10 images are visible, and only 2 of these have any indication of what company owns them. The lack of visibility or ownership attribution highlights that these images are designed to avoid detection by users and are likely a tracking tool. However, it is worth noting that even if a given image is invisible, it may be accompanied by text on the screen that tells the user that tracking is occurring. The most likely case for this happening is if the page is also setting a cookie and the user is in a European country covered by the so-called cookie law (which, globally, is not the majority of users). Based on the method used in this study, it is not possible to account for such cases should they exist. That said, half of the top 10 images belong

## Original Article

to Google, including the most requested image, the Google Analytics tracking pixel. This image is found on 46.02% of sites, is only one-by one pixel large, and is used solely for tracking purposes. The second most requested image is the Facebook “Like Button.” This image is both visible and clearly indicates the owner (albeit via iconography rather than text); its prevalence, on 21.34% of sites, demonstrates that the visual language of Facebook is having a significant influence on the Alexa top sites. It should be noted that, although this is a Facebook image, the domain hosting it is akamaihd.net. This domain belongs to the Akamai company, which provides hosting services for major brands around the world.

**Table 1. Summary of Findings.**

Rank	Top-level domain	N	% with third-party element	Average number domains contacted	% of cookie	% with Javascript
N/A	(All)	950,489	86.9	9.47	62.9	82.99
1	com	517,085	88.47	10.2	64.66	84.5
2	net	49,471	83.03	10.2	61.58	77.65
3	org	37,568	85.73	7.81	58.23	81.28
4	ru	37,560	94.33	8.48	89.42	88.72
5	de	32,447	85.34	7.52	52.87	79.26
6	uk	17,976	87.47	9.11	61.09	85.02
7	br	15,420	87.85	11	62.33	85.65
8	jp	14,478	82.05	6.87	48.08	79.52
9	pl	12,627	90.04	7.75	57.87	87.96
10	in	12,271	81.79	10.35	57.86	77.14
...						
43	edu	2,416	89.74	5.67	47.64	88.49
...						
65	gov	762	86.48	3.61	39.63	84.65



## **Original Article**

**Table 2. Top Image Characteristics.**

<b>Rank</b>	<b>Percent of sites</b>	<b>File name</b>	<b>Domain</b>	<b>Company</b>	<b>Visible</b>	<b>Indicates owner</b>
1	46.02%	_utm.gif	googleanalytics.com	Google	N	N
2	21.34%	LVx-xkvaJ0b.png	akamaihd.net	Akamai/Fac ebook	Y	Y
3	6.97%	nessie_icon_tiamat_white.png	googlesyndication.net	Google	Y	N
4	5.16%	_utm.gif	doubleclick.net	Google	N	N
5	3.81%	g.gif	wordpress.com	Wordpress	N	N
6	3.45%	vendor.gif	reson8.com	Resonate	N	N
7	2.57%	icon18_wrench_allbkg.png	blogblog.com	Google	Y	N
8	1.78%	demconf.jpg	demdex.net	Adobe	N	N
9	1.73%	small-logo.png	google.com	Google	Y	Y
10	1.69%	pixel.gif	amazonaws.com	Amazon	N	N

### **Corporate Ownership**

This article's main contribution is to reveal the corporate ownership of tracking mechanisms. This focus is vital, because the very reason that most tracking occurs is for economic motives. To understand the phenomena of Web tracking and to craft meaningful policies regulating it, it is necessary to know the actual extent and reach of the corporations that may be subject to increased regulation. The most striking finding of this study is that 78.07% of websites in the Alexa top million initiate third-party HTTP requests to a Google-owned domain. While the competitiveness of Google is well known in search, mobile phones, and display advertising, its reach in the Web

## **Original Article**

tracking arena is unparalleled. The next company, Facebook, is found on a still significant 32.42% of sites, followed by Akamai (which hosts Facebook and other companies' content) on 23.31% of sites, Twitter with 17.89%, comScore with 11.98%, Amazon with 11.72%, and AppNexus with 11.7%. These 7 companies dwarf the impact of the next 43 companies. The distribution of companies by percentage of sites tracked has a marked peak, which quickly descends into a long tail of 33 companies that track between 1% and 4% of sites.

### **Commercial Tracking and State Surveillance**

In December 2013, *The Washington Post* reported that "The National Security Agency is secretly piggybacking on the tools that enable Internet advertisers to track consumers, using 'cookies' and location data to pinpoint targets for government hacking and to bolster surveillance" (Soltani, Peterson, & Gellman, 2013, para. 1). More specifically, internal NSA documents leaked to the *Post* by former NSA contractor Edward Snowden revealed that a Google cookie named PREF was being used to track targets online. Additional documents provided to *The Guardian* by Snowden detailed that another Google cookie (DoubleClick's id), was also used by the NSA—in this case, to attempt to compromise the privacy of those using the anonymity-focused Tor network (National Security Agency, 2013).

As noted above, in the course of analyzing the Alexa top 1 million websites, 7,564,492 unique cookies were collected. Investigation of this cookie collection revealed 81,699 pages (8% of the total) spawning cookies with the name PREF set by Google-owned domains, and an additional 180,212 pages (19% of the total) spawning cookies with the name id from the (also Google-owned) DoubleClick domain. Because the Snowden disclosures have often raised more questions than answers, it is unclear whether the identified cookies are being used to surveil users today, and companies such as Google are working admirably hard to improve security. However, it is clear that the widely deployed tracking mechanisms identified in this article are of interest to more than just advertisers. In the absence of meaningful oversight, such mechanisms are ripe for abuse—not only by the United States intelligence agencies but by states around the globe.

### **Limitations**

Although the method described in this article performs very well at a large scale, it is not without its flaws, and the counts presented in this article may fall short of the actual number of third-party HTTP requests being made. These findings therefore represent a *lower bound* of the amount of third-party HTTP requests, with the actual amount likely being slightly higher. However, given the extensive prevalence of third-party HTTP requests, this constraint only serves to highlight the magnitude and scope of the findings.

Potential undercounting may be mainly attributed to the nature of the investigation itself: Third parties have no desire to be detected by users in the vast majority of cases. As previous research into mechanisms has demonstrated, advertisers go to extreme lengths to devise techniques designed to circumvent user preferences to track their activities. For example, Google recently paid out \$39.5 million in federal and state settlements over its use of techniques that circumvented cookie preferences on the Safari Web browser (Fung, 2013). Although no one went to jail, and Google admitted no fault, it is becoming clear that advertisers regularly skirt the border between violating user preferences and outright violating the law. For this reason, any census of third-party requests on the Web is bound to be subject to an undercount because those profiting from Web tracking prefer to remain in the shadows.

An additional weakness is that PhantomJS does not support Adobe Flash, Microsoft Silverlight, or other browser-based plug-ins. Requests to download such content are not detected, and if such content subsequently initiates additional requests, those are not recorded either. Historically, Flash has been a popular means of tracking on the Web. However, Flash is not supported by the vast majority of smart phones and tablets. Given the fact that industry estimates of mobile browser usage now stand at 37.12% globally, not only are Flash elements failing to reach over a third of users, it is apparent that new developments in tracking will likely not rely on Flash. It may

## **Original Article**

be too early to declare the end of Flashbased tracking, but the conclusions of this study are not significantly skewed by this methodological weakness.

A final reason for a possible undercount is the rapid rate of ingestion. Given that webXray downloads more than 10,000 pages per hour, the volume of requests from the IP address of the server running webXray will stand out as highly irregular. Many webmasters seek to limit automated scanning behavior by blocking overly active IP addresses. Therefore, it is possible that some third-party domains would have automatically blocked the IP address, yielding an undercount. This is one of many trade-offs made to achieve high speed and low cost, and the above limitations are detailed to account for the most obvious weaknesses in the technique.

A final note is that the analysis was conducted in the United States, where there are no laws that require websites to notify users of cookies, as in the European Union. Thus, elements that are wholly invisible in the United States and elsewhere may come with visual indications for EU users. No research method is perfect, and the problems mentioned here represent the type of relative uncertainty that is common in all research. None of the problems with webXray serve to significantly undermine the extensive findings detailed in this article.

### **Discussion: Do Not Track**

Of the top 100 third-party elements discovered in this study, only 7 were images; of these images, only 2 provided any indication of who the images were owned by. This means that 98% of the identified tracking elements may never alert users to their presence. Furthermore, the 2% that indicate ownership do not advertise their purpose—only their presence. Thus, if attitudes research has proven users do not want to be tracked, and such tracking is widespread and largely undetectable, how can we bring user expectations and corporate behavior in line? How can we eliminate the socially “intolerable” from the present ad-supported Web? Three main solutions are being pursued at present: industry opt-out mechanisms, browser add-ons, and the Do Not Track mechanism. The first two solutions are ineffectual to varying degrees, and the third will need the power of legislation to be respected by corporations.

Although it is nearly impossible for non-EU users to see they are being tracked, they may infer as much from the presence of advertisements in their browser. Many such advertisements come with a very small blue arrow that, when clicked, will eventually bring the user to the “online home of the Digital Advertising Alliance’s Self-Regulatory Program for Online Behavioral Advertising. On this website a user may opt out of targeted advertising. However, while most reasonable users would assume opting out means they *will not be tracked*, the DAA interprets it to mean that users “opt out from receiving interestbased advertising.” This means that users may *still be tracked*; they just will not see tailored advertisements. This practice is misleading at best and deceptive at worst. Furthermore, the opt-out mechanism requires the user to set a large array of cookies on her or his browser. Thus, if a user has been protecting his or her privacy by turning off third-party cookies, she or he must turn them back on and potentially be exposed to *more* tracking because not all trackers are members of the DAA. The industry solution is insufficient at best and intentionally misleading at worst.

The next means of protection, the use of browser add-ons, is vastly more effective at preventing tracking but comes with several problems. First, the only means a user has for knowing that such mechanisms are available is by reading news stories that highlight the add-ons or by browsing lists of popular add-ons for various browsers (privacy add-ons routinely make top-10 lists). A user must therefore actively seek out this complex remedy and spend a significant amount of time learning about how his or her browser works. For some this may be trivial, but for many it is a difficult proposition. Even assuming a user is able to install an add-on, the user is offered incomplete protection: Popular add-ons such as Ghostery and Ad-Block+ rely on blacklists of known tracking elements and fail to prevent all tracking. A Firefox add-on, RequestPolicy (Samuel & Zhang, 2009), does block all third-party requests, but at the expense of breaking the rendering mechanisms of many pages, leaving the average user with a confused mess of text on her or his screen. Although these add-ons do, indeed, offer a

## **Original Article**

significant degree of protection, they also place the burden on users and highlight the blame-the-victim mentality at the core of the issue. Furthermore, these add-ons are not available on the stock browsers included on Android phones and tablets, though recent additions to Apple iOS 9 have introduced content blocking.

Given the shortcomings of the above two solutions, it would seem that the simplest way to reduce unwanted tracking would not be setting hundreds of cookies or installing cumbersome browser add-ons, but by having a clear way for users to signal to corporations that they should not be tracked. But what is the best way to send this signal? Once again, the lowly HTTP request is the answer. It is now possible to configure all major Web browsers (Internet Explorer, Safari, Chrome, Firefox, and Opera) to send a message within the body of a request that indicates the user does not want to be tracked; when the aforementioned user-agent is sent to a server, another field, DNT, may be sent as well. DNT may be set to 1, indicating a user does not want to be tracked, or 0, indicating the user would like to be tracked; if no value is set, no conclusions about a user's preference may be drawn (Mayer & Narayanan, 2011). It would seem the problem is therefore solved: Users may now easily tell trackers to leave them alone by changing a single browser setting.

However, DNT comes with no enforcement mechanism; it is merely a polite request, and trackers may ignore it as they choose—and indeed they do. For example, comScore, the fifth most prevalent tracker, states in its privacy policy that “comScore.com does not react to [i.e., ignores] Do Not Track signals.” Investigation of the policies of the top 10 corporations with tracking elements reveals that 9 of them do not respect the DNT header. One company, Yahoo!, respected DNT for a time, but on April 30, 2014, reversed course and stated that, “As of today, web browser Do Not Track settings will no longer be enabled on Yahoo.” The *only company* in the top 10 to respect the DNT header is Twitter. According to Twitter, “When you turn on DNT in your browser, we stop collecting the information that allows us to tailor suggestions based on your recent visits to websites that have integrated our buttons or widgets.” It is important to note that Twitter follows the reasonable definition of opt-out—that is, user data is not collected—in contrast to the DAA's solution of tailored advertising not being shown. Twitter's position is admirable and represents the standard to which other companies should aspire.

Given that 90% of the top 10 tracking companies ignore DNT, it seems that the standard is a failure. However, this does not need to be the case. As mentioned earlier, Hoofnagle, Urban, and Li (2012) have determined that 60% of those surveyed supported the DNT standard. If these users were able to have their wishes given the force of law—as should be possible in democracies—the behavior of Twitter would not represent a benevolent exception but a legally enforced rule. The fact of the matter is that the technical groundwork exists, public desire exists, and all that is still needed is for lawmakers to fulfill their duty to side with citizens against unwelcome corporate practices.

## **Conclusion**

Three ideas have hopefully been established by this article: First, Westin's socially “intolerable” use of technology has come to pass; second, Senator Church's warning of an “abyss” of technologically advanced state surveillance was prophetic; and, third, there is a way out. By using a quantitative approach toward the issue of Web tracking, it has been determined that user privacy is widely compromised by many parties working together, users often have limited means to detect these privacy breeches, a handful of powerful corporations receive the vast majority of user data, and state intelligence agencies such as the NSA may be leveraging this commercial surveillance infrastructure for their own purposes. Furthermore, current privacy protections are wholly inadequate in light of the scale and scope of the problem. Despite this, by having singled out the corporations that covertly observe users, it is now possible to identify the degree to which these companies abide by expressed user preferences, or fail to do so. A full exploration of specific policy implementations is beyond the purview of this research, but the measures contained herein may be used by the policy community to inform, support, and advance solutions that ensure that DNT moves from a polite request to an enforceable command.

## Original Article

### References

- Acar, G., Juarez, M., Nikiforakis, N., Diaz, C., Gürses, S., Piessens, F., & Preneel, B. (2013). FPdetective: Dusting the Web for fingerprinters. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer and Communications Security* (pp. 1129–1140). New York, NY: Association for Computing Machinery.
- Ackerman, M. S., Cranor, L. F., & Reagle, J. (1999). Privacy in e-commerce: Examining user scenarios and privacy preferences. In *Proceedings of the 1st ACM Conference on Electronic Commerce* (pp. 1–8). New York, NY: Association for Computing Machinery.
- Ayenson, M., Wambach, D., Soltani, A., Good, N., & Hoofnagle, C. (2011, July 29). *Flash cookies and privacy II: Now with HTML5 and eTag respawning*. Retrieved from [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=1898390](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1898390)
- Bamford, J. (2011, September 8). Post-9/11, NSA “enemies” include us. *Politico*. Retrieved from <http://www.politico.com/story/2011/09/post-9-11-nsa-enemies-include-us-062999>
- Castelluccia, C., Grumbach, S., & Olejnik, L. (2013). Data harvesting 2.0: From the visible to the invisible Web. In *Proceedings of the Twelfth Workshop on the Economics of Information Security (WEIS 2013)* (pp. 1–14). Washington, DC: Georgetown University.
- Eckersley, P. (2010, January). How unique is your Web browser? In *Proceedings of the Privacy Enhancing Technologies Symposium* (pp. 1–18). Berlin, Germany: Springer.
- Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., & Berners-Lee, T. (1999). *Hypertext transfer protocol—HTTP/1.1*. Internet Engineering Task Force. Retrieved from <http://www.ietf.org/rfc/rfc2616.txt>
- Fung, B. (2013, November 19). Why states are the big winner in the \$17 million Google-Safari settlement. *The Washington Post*. Retrieved from <https://www.washingtonpost.com/news/theswitch/wp/2013/11/19/why-states-are-the-big-winner-in-the-17-million-google-safarisettlement/>
- Hoofnagle, C. J., Urban, J. M., & Li, S. (2012, October). *Privacy and modern advertising: Most US Internet users want “do not track” to stop collection of data about their online activities*. Retrieved from [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2152135&rec=1&srcabs=2103405&alg=1&pos=2](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2152135&rec=1&srcabs=2103405&alg=1&pos=2)
- Jackson, C., Bortz, A., Boneh, D., & Mitchell, J. C. (2006). Protecting browser state from Web privacy attacks. In *Proceedings of the 15th International Conference on World Wide Web* (pp. 737–744). New York, NY: Association for Computing Machinery.
- Jang, D., Jhala, R., Lerner, S., & Shacham, H. (2010). An empirical study of privacy-violating information flows in Javascript Web applications. In *Proceedings of the 17th ACM Conference on Computer and Communications Security* (pp. 270–283). New York, NY: Association for Computing Machinery.



## Original Article

- Krishnamurthy, B., & Wills, C. E. (2006). Generating a privacy footprint on the Internet. In *Proceedings of the 6th ACM SIGCOMM Conference on Internet Measurement*, (pp. 65–70). New York, NY: Association for Computing Machinery.
- Krishnamurthy, B., & Wills, C. E. (2009). Privacy diffusion on the Web: A longitudinal perspective. In *Proceedings of the 18th International Conference on World Wide Web* (pp. 541–550). New York, NY: Association for Computing Machinery.
- Kristol, D., & Montulli, L. (1997). *HTTP state management mechanism*. Internet Engineering Task Force. Retrieved from <https://www.rfc-editor.org/rfc/rfc2109.txt>
- Libert, T. (2015). Privacy implications of health information seeking on the Web. *Communications of the ACM*, 58(3), 68–77.
- Madden, M., Rainie, L., Zickuhr, K., Duggan, M., & Smith, A. (2014). *Public perceptions of privacy and security in the post-Snowden era*. Washington, DC: Pew Research Center. Retrieved from [http://www.pewinternet.org/files/2014/11/PI\\_PublicPerceptionsofPrivacy\\_111214.pdf](http://www.pewinternet.org/files/2014/11/PI_PublicPerceptionsofPrivacy_111214.pdf)
- Mayer, J. R., & Mitchell, J. C. (2012). Third-party Web tracking: Policy and technology. In *Proceedings of the 2012 IEEE Symposium on Security and Privacy* (pp. 413–427). New York, NY: Institute of Electrical and Electronics Engineers.
- Mayer, J. R., & Narayanan, A. (2011). *Do not track: A universal Third-party Web tracking opt-out*. Internet Engineering Task Force. Retrieved from <http://tools.ietf.org/html/draft-mayer-do-not-track-00>
- National Security Agency. (2013). “Tor stinks” presentation. Retrieved from <http://www.theguardian.com/world/interactive/2013/oct/04/tor-stinks-nsa-presentationdocument>
- Nikiforakis, N., Kapravelos, A., Joosen, W., Kruegel, C., Piessens, F., & Vigna, G. (2013, May). Cookieless monster: Exploring the ecosystem of Web-based device fingerprinting. In *Proceedings of the Security and Privacy (SP), 2013 IEEE Symposium* (pp. 541–555). New York, NY: Institute of Electrical and Electronics Engineers.
- Roesner, F., Kohno, T., & Wetherall, D. (2012). Detecting and defending against third-party tracking on the Web. In *Proceedings of the 9th USENIX Conference on Networked Systems Design and Implementation* (pp. 12–12). Berkeley, CA: USENIX Association.
- Samuel, J., & Zhang, B. (2009). RequestPolicy: Increasing Web browsing privacy through control of crosssite requests. In *Proceedings of the Privacy Enhancing Technologies Symposium* (pp. 128–142). Berlin, Germany: Springer.
- Soltani, A., Peterson, A., & Gellman, B. (2013, December 10). NSA uses Google cookies to pinpoint targets for hacking. *The Washington Post*. Retrieved from <https://www.washingtonpost.com/blogs/the-switch/wp/2013/12/10/nsa-uses-google-cookies-topinpoint-targets-for-hacking>

**Original Article**

- Turow, J. (2003). *Americans and online privacy: The system is broken*. Philadelphia, PA: Annenberg Public Policy Center, University of Pennsylvania.
- Turow, J. (2011). *The daily you: How the new advertising industry is defining your identity and your worth*. New Haven, CT: Yale University Press.
- Turow, J., Hennessy, M., & Draper, N. A. (2015). *The tradeoff fallacy: How marketers are misrepresenting American consumers and opening them up to exploitation*. Philadelphia, PA: Annenberg School for Communication, University of Pennsylvania.
- Turow, J., King, J., Hoofnagle, C. J., Bleakley, A., & Hennessy, M. (2009). *Americans reject tailored advertising and three activities that enable it*. Retrieved from [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=1478214](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1478214)
- Westin, A. (1967). *Privacy and freedom*. London, UK: Bodley Head.
- Yen, T.-F., Xie, Y., Yu, F., Yu, R. P., & Abadi, M. (2012). Host fingerprinting and tracking on the Web: privacy and security implications. In *Proceedings of Network and Distributed System Security Symposium* (pp. 1–16). Reston, VA: Internet Society.