Empowering Users to Express a "Do Not Track" Rule: A Step Toward Conveying User Privacy Preferences

W3C Web Tracking and User Privacy Workshop
April 28-29, 2011

John Morris and Alissa Cooper
Center for Democracy & Technology

1. Introduction

This paper considers the privacy protection theory that underlies the "do not track" header and DOM property proposals, and addresses (and tries to rebut) a number of arguments that have been advanced against similar proposals considered elsewhere within the W3C.[1] Fundamentally, the idea of these do not track proposals is that a user-set privacy rule – an instruction not to track the user's browsing – is included as a header in HTTP requests transmitted to web servers and/or as a property in the DOM. These approaches are instances of a broader approach to privacy protection – that of conveying user privacy preferences to entities that are in position to act on them

We believe that conveying user privacy preferences in general – and the do not track header and DOM property proposals in particular – have strong potential to be of value in the effort to protect privacy. There are other possible approaches to achieve a do not track regime, as summarized in a companion submitted paper entitled "Summary Comparison of Universal Opt-Out Mechanisms for Web Tracking" (and as detailed more fully in an Internet Draft recently submitted to the IETF[2]). The authors of both submitted papers are supportive of all of the approaches, and indeed we believe that a number of complimentary approaches could be implemented. We focus on the header and DOM property here for the purposes of drawing comparisons to previous efforts in the W3C, and urging further consideration of more broadly allowing users to set and convey their privacy rules.

The general approach of conveying user privacy preferences has been considered in at least three separate contexts within Internet standards bodies. First, starting in 2001, the Geopriv Working Group at the IETF implemented this approach by attaching privacy rules to location data.[3] Second, the Geolocation WG of the W3C considered a similar approach,

---

[1] Many of the points made in this paper were first articulated in a paper, "Binding Privacy Rules to Data: Empowering Users on the Web," submitted by the authors (and Erica Newland) to the W3C Privacy Workshop held in July, 2010. That paper focused more broadly than just on "do not track" proposals, which are instances of the broader concept of allowing users to set and transmit rules to restrict third party use of their information and activities.

[2] *See* "Overview of Universal Opt-Out Mechanisms for Web Tracking," available at http://tools.ietf.org/html/draft-cooper-web-tracking-opt-outs-00.

[3] *See* http://datatracker.ietf.org/wg/geopriv/charter/. One of the authors of this paper, Alissa Cooper, is a co-chair of the Geopriv WG, and the other, John Morris, is a co-author of a number of Geopriv RFCs.

but rejected it for many of the reasons discussed (and rebutted) in this paper.[4] Third, the Device API and Policy Working Group (DAP) of the W3C has explored the notion of passing user "privacy rulesets" to consumers of the DAP APIs,[5] but the proposal has met resistance from some WG participants for many of the reasons discussed in this paper.[6] This paper briefly addresses the criticisms raised by opponents to the approach of conveying user privacy preferences, and assesses those criticisms in the context of the do not track header and DOM property proposals.

2. Conveying User Privacy Preferences

The central feature of conveying user privacy preferences is that when a user communicates with another entity, applicable privacy rules are also conveyed to the entity to ensure that entities that receive information about the user are informed of how they may (or may not) use it. By creating a structure to convey the users' preferences along with their information or communications, the likelihood that those preferences will be honored necessarily increases. In particular, no entity can disavow knowledge of users' preferences for how their information may or may not be used. Conveying user privacy preferences allows users to express their desire for and expectations of privacy, which in turn helps to bolster social and legal systems' protection of those expectations.

Applying and affixing usage rules to information is a well-known way of protecting information, long before the World Wide Web (for example, by placing the © copyright symbol on documents). More recently, the Creative Commons[7] model is one prominent example, allowing an owner of a work to set four types of rules ("Attribution," "Noncommercial," "No Derivative Works" and "ShareAlike") governing the subsequent use of the work. After the author sets these rules, the rules are conveyed together with the work itself, so that every consumer of the work is aware of the copyright terms.

Another example where usage rules are bound to data is in security classification systems (such as marking documents with a "Secret" designation). As these examples reveal, these systems of rule enforcement are *not* self-executing. Unlike some technical strategies (such as encryption), these systems rely on external, *non-technical* mechanisms (such as laws, contracts, or company rules) to enforce the protection of the information. The do not track header and DOM property proposals follow this model – they propose the creation of a technical requirement to ensure that the applicable user preference is always conveyed to entities capable of tracking, and it leaves to regulatory, legal, and market forces the enforcement of the user's directive.

---

[4] *See* http://www.w3.org/2008/geolocation/.
[5] *See* http://dev.w3.org/2009/dap/privacy-rulesets/.
[6] *See* http://www.w3.org/2009/dap/. The authors of this paper have been among the primary advocates for the approach of binding user-set rules to data within both the GeoLocation and DAP working groups.
[7] *See* http://creativecommons.org/.

3. Arguments Against Conveying User Privacy Preferences

In the Geolocation WG and, to a lesser extent, the DAP WG, a number of arguments have been raised against the idea of conveying user privacy preferences. This section briefly recaps some of the criticisms and responses to them, without intending to be an exhaustive discussion of either side of the arguments.

a. Conveying user privacy preferences does not protect privacy through technical means (such as encryption).

Conveying preferences does not, by itself, provide technical means through which it can be reasonably guaranteed that users' privacy rules will be honored by recipients of their data. Thus, the transmission of a do not track header, for example, does not in any technical way assure that the recipient web site receiving the header will not, in fact, track the user. Instead, the privacy protection is provided by virtue of the fact that data recipients are informed of the user's preference, and are expected to only use data in accordance with that preference.

By conveying the user's preference, the approach provides valuable information so that *non-technical* forces such as legal contracts, governmental consumer protection authorities, and marketplace feedback can better enforce those preferences. If a commercial recipient violates a user's clear privacy preference, for example, the recipient can, in a growing number of countries, be charged with violating consumer or data protection laws. In the absence of an expressed preference, consumer protection authorities are less able to protect consumers whose information has been abused.

b. Implementing a preference interface in a user agent would be hard, and users might be confused.

Without question user interfaces are hard. But given that the user agent serves as a crucial gateway between users and the web, providing centralized privacy preference interfaces in the user agent may in fact help to reduce the confusion caused by each individual web site or app giving users different controls and interfaces over essentially the same user data being communicated through the user agent. User agents already contain privacy preference interfaces, for example to control cookies. When cookies were first introduced on the web, browsers provided no way for users to control their use.[8] As concerns were raised about potentially privacy-invasive uses of cookies, browser vendors began to add cookie controls into their products, beginning with rudimentary tools and evolving over time to the more sophisticated controls in place today. Browser makers continue to explore simple ways to present privacy choices in the browser,[9] and interfaces to convey user preferences should be part of that exploration.

---

[8] *See Federal Trade Commission Staff Report. Public Workshop on Consumer Privacy on the Global Information Infrastructure, Part III: Enhancing Consumer Protection Online* (Dec. 1996), available at http://www.ftc.gov/reports/privacy/Privacy4.shtm.

[9] *See, e.g.,* http://people.mozilla.com/~faaborg/files/firefox4Mockups/prefPaneWebSites-i2.png.

c. Users would blame the browser when web sites violate the users' expressed preferences.

If a browser provides a user interface allowing users to set a privacy preference (such as do not track) and those rules are later violated, there is a risk that the browser will be blamed. There are, however, affirmative steps that can be taken when designing the user interface to mitigate this possibility. A user interface can make clear that it is soliciting preferences to be conveyed to the recipient of the information, and that the recipient is responsible for honoring them. The user interface associated with the do not track header in Firefox 4 provides a good example: users can check a box to "Tell web sites I do not want to be tracked," as opposed to a box that says "Do not let web sites track me."[10] By being careful to convey the limits of the browsers' control over later tracking or other uses of the users' data and providing supplemental user education about the user agent privacy settings,[11] the user agent can reduce the risk that it would be blamed for a privacy violation by a receiving entity.

d. Rather than providing incomplete privacy protection, it is better for users to think there is no privacy protection.

In the security context, there may be real risk if users mistake weak protection for adequate protection – they may expose critical data (such as, say, bank account login information) and then suffer catastrophic harm. And there often is an available way to achieve real security, even if it means a delay or inconvenience in performing a transaction.

The privacy context is quite different. The harm is often more incremental, and users are better off if even a subset of recipients of their information or communications honor their privacy preferences. In the web tracking context, users may not fully understand the extent to which data about their web behavior is being collected and used, but they may still want to convey their preference not to be tracked, even if it is not universally honored from the outset. In the context of web applications that require users to affirmatively share data about themselves, users are often presented with a "Hobson's choice" with regards to their data: using a service requires implicit acceptance of all future data uses by the service provider, and the only other option is to not use the service at all. Unlike in the security context, users often have no alternative to this "take-it-or-leave-it" approach to privacy, and so users are forced to give up their privacy. Any enhanced privacy protections, even if incomplete, will offer users a substantive improvement over the status quo.

f. We are not sure it will work.

The approach of sending user privacy preferences is new to the applications layer, and there is certainly no guarantee that this framework will succeed. But, one thing is certain:

---

[10] *See* http://support.mozilla.com/en-US/kb/how-do-i-stop-websites-tracking-me.
[11] *See, e.g.,* http://ie.microsoft.com/testdrive/Browser/TrackingProtection/Default.html, http://support.mozilla.com/en-US/kb/how-do-i-stop-websites-tracking-me.

the status quo has failed to provide meaningful privacy protection on the web. Privacy policies are not widely read or understood[12] while web tracking continues to become more sophisticated and pervasive[13] despite users' discomfort with it.[14] Doing nothing to change this situation in the face of the growth of web applications will only further jeopardize user privacy on the web.[15]

* * * *

By using a do not track header and/or DOM property, users can be given some element of control, as well as some legal claim, over whether their web browsing is tracked.[16] The same argument could be made for building mechanisms that would allow users to express their privacy preferences over geolocation and other types of sensitive personal information. The work of some browser vendors to implement a do not track header and/or DOM property is strong evidence that the objections to the approach of conveying user privacy preferences are in fact surmountable. As with do not track, placing users in the position of being able to set – and have expressed in a standardized way – their privacy preferences will greatly increase the chance that those preferences are honored.

[12] *See, e.g.*, http://portal.acm.org/citation.cfm?id=1614511.
[13] *See, e.g.*, http://web.cs.wpi.edu/~cew/papers/soups07.pdf.
[14] *See, e.g.*, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1478214.
[15] Significantly, legislators and regulators in the United States are increasingly interested in the potential that do not track techniques offer. Although the U.S. Federal Trade Commission has authority to regulate unfair and deceptive trade practices – and could pursue violations of do not track instructions based on its current authority – some in Congress are considering proposals to make explicit the FTC's authority to enforce do not track rules set by users.
[16] As noted above, the authors do not believe that a do not track header and/or DOM property are the only approaches to web tracking opt outs that are worth considering, and other approaches may also be useful tools to provide to users.