Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 8, 2011

J. Mayer
A. Narayanan
Stanford University
S. Stamm
Mozilla
March 7, 2011

# Do Not Track: A Universal Third-Party Web Tracking Opt Out
draft-mayer-do-not-track-00

## Abstract

This document defines the syntax and semantics of Do Not Track, an HTTP header-based mechanism that enables users to express preferences about third-party web tracking. It also provides a standard for how web services should comply with such user preferences.

## Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 8, 2011.

## Copyright Notice

Table of Contents

## 1. Recognition

The Do Not Track effort is much broader than this standards document, and we recognize the following individuals without whom Do Not Track would not be possible. For a detailed history of Do Not Track, see [HistoryOfDNT]. We particularly laud the efforts of Christopher Soghoian, whose tireless advocacy led Do Not Track from a technical prototype to a leading privacy proposal.

### 1.1. Contributors

Alissa Cooper
Center for Democracy and Technology

Christopher Soghoian
Indiana University

Ashkan Soltani

Harlan Yu
Princeton University

### 1.2. Acknowledgements

Peter Eckersley
Electronic Frontier Foundation

Alexander Fowler
Mozilla

John Mitchell

Stanford University

Lee Tien
Electronic Frontier Foundation

## 2. Introduction

The content of a website is increasingly sourced from numerous entities. This development has given many companies the ability to track users across millions of sites. A number of services now exist





solely to track users, often via invisible embedded content. Users widely perceive such third-party tracking as an invasion of privacy (see [WebsNewGoldMine] and [Turow09]).

The explosion of stateful (see [Evercookie], [Aggarwal10], and [McKinley08]) and stateless (see [Eckersley10] and [Mayer09]) techniques for tracking users, accompanied by the proliferation of third-party tracking (see [Krishnamurthy10]), prohibit a purely technical means of preventing tracking. Do Not Track is instead a means of allowing users to express their preferences about tracking, including to opt out of tracking some or all of the time.

A preference signaling mechanism can, of course, be ignored by bad actors. But the most pervasive third-party trackers are law-abiding commercial enterprises (see [Krishnamurthy10]). This standard intends to aid these fair players by allowing them to honor a user's preferences.

## 3. Definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

This specification uses the Augmented Backus-Naur Form (ABNF) notation of [RFC5234].

The terms user agent, server, proxy, header, request, and response

have the same meaning as in the HTTP/1.1 specification ([RFC2616]).

"Explicit user consent" means a user is likely to understand and accept the choice she makes. Agreement to a terms of service or privacy policy does not, in general, constitute explicit user consent.

A "functional entity" is a commercial, nonprofit, or government organization, a subsidiary or unit of such an organization, or a person.

"THIRD-PARTY TRACKING" is shorthand for activities covered by Section 9.1 and Section 9.2, and not excepted by Section 9.3.

A "public suffix" is a domain name under which users can register domain names. A list is maintained at [PublicSuffix]. This document uses public suffixes instead of top-level domains (see [RFC0920]) because they more accurately reflect organizational boundaries.





"Protocol logs" means logs for all network protocols that arise from an HTTP request and response.

## 4. Overview

This document is organized into two parts. The first part details the technical implementation of Do Not Track: the header syntax (Section 5), user agent requirements (Section 6), intermediary requirements (Section 7), and server requirements (Section 8). The second part provides the policy a server implementing Do Not Track must observe (Section 9).

### 4.1. Example

In the status quo: A user navigates a sequence of popular websites, many of which incorporate content from a major advertising network. In addition to delivering advertisements, the advertising network assigns a unique cookie to the user agent and compiles observations of the user's browsing habits.

With Do Not Track: A user enables Do Not Track in her web browser.

She navigates a sequence of popular websites, many of which incorporate content from a major advertising network. The advertising network delivers advertisements, but refrains from THIRD-PARTY TRACKING of the user.

## 5. Header Syntax

The Do Not Track HTTP header, "DNT", must take one of two values: "1" ("opt out") or "0" ("opt in"). All other values are reserved.

DNT = "DNT" ":" BIT

For clarity this document refers to an opt-out header as OPT-OUT, an opt-in header as OPT-IN, and the absence of a header as NO-EXPRESSED-PREFERENCE.

## 6. User Agent Requirements

### 6.1. OPTIONAL Support

A user agent MAY include a Do Not Track header in any HTTP request.





### 6.2. User Interface RECOMMENDED

A user agent that implements Do Not Track SHOULD provide a user interface for modifying preferences. The user interface design is left to the user agent.

### 6.3. Default

A user agent MAY adopt NO-EXPRESSED-PREFERENCE or OPT-OUT by default. It MUST NOT transmit OPT-IN without explicit user consent.

## 7. Intermediary Requirements

A proxy or other intermediary MUST NOT add, remove, or modify a Do

Not Track header without explicit user consent.

## 8. Server Requirements

### 8.1. Opt Out

In processing a request that includes an OPT-OUT header, a server MUST NOT perform THIRD-PARTY TRACKING. The server MUST instruct the user agent to delete any data previously stored for THIRD-PARTY TRACKING.

### 8.2. Opt In

In processing a request that includes an OPT-IN header, a server MAY perform THIRD-PARTY TRACKING.

### 8.3. Header Not Present

In processing a NO-EXPRESSED-PREFERENCE request, a server MAY perform THIRD-PARTY TRACKING. The functional entity responsible for the server MUST NOT draw any inferences about a user's preferences from the absence of an OPT-OUT or OPT-IN header.

### 8.4. Response Header RECOMMENDED

In responding to a request that includes a Do Not Track header, a third-party server that complies with Do Not Track SHOULD echo the request header. For example:

```
GET /thirdpartycontent.html HTTP/1.1
Host: thirdparty.example.com
```





```
DNT: 1

HTTP/1.1 200 OK
Date: Mon, 7 March 2011 01:23:45 GMT
Server: Apache/2.2.17 (Unix)
Content-Length: 123
Connection: close
Content-Type: text/html; charset=UTF-8
```

DNT: 1

This feature is intended to aid in the decentralized collection of statistics about the Do Not Track mechanism, including adoption rates and intermediary operations. It is also intended to clearly identify whether a request was processed in compliance with Do Not Track.

## 9. Server Policy

This section specifies the requirements for server compliance with a Do Not Track OPT-OUT: A server acting in a third-party capacity (see Section 9.1) MUST NOT track (see Section 9.2) a user or user agent unless subject to an exception (see Section 9.3).

### 9.1. Definitions of "First Party" and "Third Party"

A first party is a functional entity with which the user reasonably expects to exchange data. In most cases the functional entity responsible for the web page a user has navigated to is the sole first party.

A third party is a functional entity with which the user does not reasonably expect to share data. In general advertising networks, analytics services, and social plug-in providers are third parties. To a first approximation, a functional entity is a third party if it differs from the current page in:

1. Public suffix plus one domain name (PS+1), or
2. PS+1 authoritative name servers, or
3. PS+1 of CNAME records.

We emphasize that this rule is only an approximation. Many first parties span several domain names, and many third parties are located at a subdomain of a first party.

In practice a third party usually interacts with a user agent via content embedded on a first-party webpage. A third party could also receive data from a first party.





### 9.2. Definition of "Tracking"

Tracking includes collection, retention, and use of all data related to the request and response.

### 9.3. Exceptions

As a general guideline, exceptions to Do Not Track are warranted when commercial interests substantially outweigh privacy and verification interests. The following activities are excepted:

1. Tracking of users who have explicitly consented to tracking, such as by enabling a checkbox in a preferences menu on the first-party website of the tracking service.
2. Data obtained by a third party exclusively on behalf of and for the use of a first party.
3. Data that is, with high confidence, not linkable to a specific user or user agent. This exception includes statistical aggregates of protocol logs, such as pageview statistics, so long as the aggregator takes reasonable steps to ensure the data does not reveal information about individual users, user agents, devices, or log records. It also includes highly non-unique data stored in the user agent, such as cookies used for advertising frequency capping or sequencing. This exception does not include anonymized data, which recent work has shown to be often re-identifiable (see [Narayanan09] and [Narayanan08]).
4. Protocol logs, not aggregated across first parties, and subject to a two week retention period.
5. Protocol logs used solely for advertising fraud detection, and subject to a one month retention period.
6. Protocol logs used solely for security purposes such as intrusion detection and forensics, and subject to a six month retention period.
7. Protocol logs used solely for financial fraud detection, and subject to a six month retention period.

To ensure data allowed for only specific uses is adequately protected, functional entities SHOULD implement strong internal controls.

## 10. Implementation Considerations

### 10.1. Selective Opt Out and Opt In RECOMMENDED

A user agent implementing Do Not Track SHOULD allow a user to selectively opt out of or opt into tracking on specific first-party websites, by specific third parties, or by specific third parties on

specific first-party websites. Definition and implementation of selective opt out and opt in is outside the scope of this document.

### 10.2. Verification

Verification systems may be needed to ensure compliance with Do Not Track. Such systems are outside the scope of this document.

## 11. Security Considerations

This document does not introduce any known security considerations.

## 12. Privacy Considerations

User agent implementation of Do Not Track contributes a small amount of fingerprintable information (see [Eckersley10] and [Mayer09]). The amount of information depends on the degree of adoption. Supposing, for example, that 10% of user agents have Do Not Track enabled, the header adds only -log2(0.1) (roughly 3.3) bits of identifying information to the user agent. Relative to other sources of fingerprintable information Do Not Track is of minimal concern.

## 13. IANA Considerations

This specification calls for a new IANA provisional message header field registration, in accordance with [RFC3864].

Header field name: see Section 5
Applicable protocol: http ([RFC2616])
Status: standard
Author/Change controller: IETF
Specification document: this document

## 14. References

### 14.1. Normative References

[RFC0920] Postel, J. and J. Reynolds, "Domain requirements", RFC 920, October 1984.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate

Requirement Levels", BCP 14, RFC 2119, March 1997.





[RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.

[RFC3864] Klyne, G., Nottingham, M., and J. Mogul, "Registration Procedures for Message Header Fields", BCP 90, RFC 3864, September 2004.

[RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.

## 14.2. Informative References

[Aggarwal10] Aggarwal, G., Bursztein, E., Jackson, C., and D. Boneh, "An Analysis of Private Browsing Modes in Modern Browsers", 2010, <http://crypto.stanford.edu/~dabo/pubs/papers/privatebrowsing.pdf>.

[Eckersley10] Eckersley, P., "How Unique Is Your Web Browser?", 2010, <https://panopticlick.eff.org/browser-uniqueness.pdf>.

[Evercookie] Kamkar, S., "Evercookie", September 2010, <http://samy.pl/evercookie/>.

[HistoryOfDNT] Soghoian, C., "The History of the Do Not Track Header", January 2011, <http://paranoia.dubfire.net/2011/01/history-of-do-not-track-header.html>.

[Krishnamurthy10] Krishnamurthy, B., "Privacy Leakage on the Internet", March 2010, <http://www.ietf.org/proceedings/77/slides/plenaryt-5.pdf>.

[Mayer09] Mayer, J., ""Any person... a pamphleteer": Internet

Anonymity in the Age of Web 2.0", April 2009,
<http://stanford.edu/~jmayer/papers/thesis09.pdf>.

[McKinley08]
McKinley, K., "Cleaning Up After Cookies", December 2010,
<https://www.isecpartners.com/files/
iSEC_Cleaning_Up_After_Cookies.pdf>.

[Narayanan08]





Narayanan, A. and V. Shmatikov, "Robust De-anonymization
of Large Sparse Datasets", 2008,
<http://www.cs.utexas.edu/~shmat/shmat_oak08netflix.pdf>.

[Narayanan09]
Narayanan, A. and V. Shmatikov, "De-anonymizing Social
Networks", 2009,
<http://www.cs.utexas.edu/~shmat/shmat_oak09.pdf>.

[PublicSuffix]
"The Public Suffix List", <http://publicsuffix.org/>.

[Turow09] Turow, J., King, J., Hoofnagle, C., Bleakley, A., and M.
Hennessy, "Americans Reject Tailored Advertising and Three
Activities that Enable It", September 2009,
<http://ssrn.com/abstract=1478214>.

[WebsNewGoldMine]
Angwin, J., "The Web's New Gold Mine: Your Secrets",
July 2010, <http://online.wsj.com/article/
SB10001424052748703940904575395073512989404.html>.

## Authors' Addresses

Jonathan Mayer
Stanford University

URI: http://jonathanmayer.net

Arvind Narayanan

Stanford University

URI: http://randomwalker.info

Sid Stamm
Mozilla

URI: http://sidstamm.com