Updated Proposal - Outline in preparation for presentation in Seattle

From: Shane Wiley <wileys@yahoo-inc.com>
Date: Mon, 11 Jun 2012 19:25:41 -0700
To: "public-tracking@w3.org" <public-tracking@w3.org>
Message-ID: <63294A1959410048A33AEE161379C8023D18786247@SP2-EX07VS02.ds.corp.yahoo.com>

Hello TPWG,

Due to "recent activities" I'm a bit behind on providing the final presentation for our updated proposal in preparation for Seattle. We'll be reviewing this in more detail in Seattle but I wanted to share some of the initial elements up-front so we have time as a working group to begin discussion and consider perspectives leading up to the meeting.

Goal: Evolve DC proposal to bridge the divide with the advocate proposal and set a final recommendation for these elements

- Definition of First Party

o Advocate Position: Common Branding

o Industry Position: Affiliate

o Concession Proposal: Affiliate with "easy discoverability" ("Affiliate List" within one click from each page or owner clearly identified within one click from each page. For example, a link in the privacy policy would meet this requirement.)

- Permitted Uses

o Advocate Position: Unlinkable Data w/ arbitrary "grace period"

o Industry Position: Enumerated uses, broadly scoped, general data minimization

o Concession Proposal: Tightened up permitted uses, narrowly and strictly scoped, data minimization focus with required transparency, reasonable safeguards, defined unlinkable (highlighting this moves resulting data outside of scope)

- For All Permitted Uses

o What won't occur: Outside of Security, all other permitted uses will not allow for altering a specific user's online experience (no profiling, no further alteration to the user experience base on profiled information)

o Data Minimization: Each organization engaging in Permitted Uses and claiming W3C DNT compliance, must provide public transparency of their data retention period (may enumerate each individually if they vary across Permitted Uses)

o Reasonable Safeguards: Reasonable technical and organizational safeguards to

prevent further processing: collection limitations, data siloing, authorization restrictions, k-anonymity, unlinkability, retention time, anonymization, pseudonymization, and/or data encryption.

· Permitted Uses: Security/Fraud, Financial Logging/Auditing, Frequency Capping, Debugging, Aggregate Reporting*

o For each Permitted Use:

§ (Normative) Detailed, singular business purpose description

§ (Non-normative) Will explain why the processing with identifiers is proportionate
*NOTE - Aggregate Reporting covers general analytics needs, product improvement, and market research uses

· Explicit and Separate User Choice

o User must expressly activate DNT signal (TPWG already agreed on this point)

o Servers may respond to users that their UA is "invalid" if they believe this to be the case (on the hook to defend this position)

o Efforts to misled users to activate DNT will be seen as "invalid"

· With this Proposal

o Users gain a consistent, local tool to communicate their opt-out preference (avoids property specific opt-out pages)

o The users choice is persistent for each device/UA (avoids accidental deletion)

o Outside of Security purposes, the user will no longer experience alterations to their online experiences derived from multi-site activity

o Only minimal data is retained for necessary business operations and retention periods are transparent to users

o All "harms" are removed (outside of government intrusion risk where there are no documented cases of this occurring with 3rd party anonymous log file data)

· Unlinkability

<Normative>

Un-linkable Data is outside of the scope of the Tracking Preference standard as information is no longer reasonably linked to a particular user, user agent, or device.

Definition: A dataset is un-linkable when reasonable steps have been taken to modify

data such that there is confidence that it contains only information which could not be linked to a particular user, user agent, or device.

<Non-Normative>

There are many valid and technically appropriate methods to de-identify or render a data set "un-linkable". In all cases, there should be confidence the information is unable to be reverse engineering back to a "linkable" state. Many tests could be applied to help determine the confidence level of the un-linking process. For example, a k-anonymous test could be leveraged to determine if the mean population resulting from a de-linking exercise meets an appropriate threshold (a high-bar k-anonymous threshold would be 1024).

As there are many possible tests, it is recommended that companies publically stating W3C Tracking Preference compliance provide transparency to their delinking process so external experts and auditors can assess if they feel this steps are reasonable given the risk of a particular dataset.

· Information That Is Un-linkable When Collected: A third party may collect non-protocol information if it is, independent of protocol information, un-linkable data. The data may be retained and used subject to the same limitations as protocol information.

Example: Example Advertising sets a language preference cookie that takes on few values and is shared by many users.

· Information That Is Un-linkable After Aggregation: During the period in which a third party may use protocol information for any purpose, it may aggregate protocol information and un-linkable data into an un-linkable dataset. Such a dataset may be retained indefinitely and used for any purpose.

Example: Example Advertising maintains a dataset of how many times per week Italy-based users load an ad on Example News.

· Information That Is Un-linkable After Anonymization: At some point after collection, a unique ID from a product cookie has a one-way salted hash applied to the identifier to break any connection between the resulting dataset and production identifiers. To further remove dictionary attacks on this method, its recommended that "keys" are rotated on a regular basis.

Received on Tuesday, 12 June 2012 02:26:41 GMT

Do Not Track — Compromise Proposal

Unofficial Draft 06 June 2012

Editors:
Peter Eckersley, Electronic Frontier Foundation
Tom Lowenthal, Mozilla
Jonathan Mayer, Stanford University

Abstract

Abstract, version, and status information are not relevant in this partial draft.

Status of This Document

This document is merely a public working draft of a potential specification. It has no official standing of any kind and does not represent the support or consensus of any standards organisation.

Table of Contents

A.2 Informative references

1. User Agents

1.1 Explicit Consent Requirement

Note: This section was recently added and has not been extensively discussed with stakeholders. Please consider it a preliminary position.

An ordinary user agent MUST NOT send a Tracking Preference signal without a user's explicit consent.

Example: The user agent's privacy preferences pane includes controls for configuring the Tracking Preference signal.

Example: On first run, the user agent prompts the user to configure the Tracking Preference signal.

2. Parties, First Parties and Third Parties

2.1 Parties

2.1.1 Definitions

A ***functional entity*** is any commercial, nonprofit, or governmental organization, a subsidiary or unit of such an organization, or a person.

Functional entities are ***affiliated*** when they are related by both common majority ownership and common control.

A ***party*** is a set of functional entities that are affiliated.

2.1.2 Transparency

2.1.2.1 Requirement

A functional entity must make its affiliated functional entities easily discoverable by a user.

2.1.2.2 Non-Normative Discussion

Affiliation may be made easily discoverable by a user in many ways, including but not limited to: prominent and common branding on pages, one click away within a privacy policy, or a machine-readable format in a well-known location.

2.2 Network Interaction

2.2.1 Definition

A ***network interaction*** is an HTTP request and response, or any other set of logically related network traffic.

2.2.2 Non-Normative Discussion

Determination of a party's status is limited to a single transaction because a party's status may be affected by time, context, or any other factor that influences user expectations.

2.3 First Parties and Third Parties

2.3.1 Definitions

A ***first party*** is any party, in a specific network interaction, that can infer with high probability that the user knowingly and intentionally communicated with it. Otherwise, a party is a third party.

A ***third party*** is any party, in a specific network interaction, that cannot infer with high probability that the user knowingly and intentionally communicated with it.

2.3.2 Non-Normative Discussion

2.3.2.1 Overview

We draw a distinction between those parties an ordinary user would or would not expect to share information with, "first parties" and "third parties" respectively. The delineation exists for three reasons.

First, when a user expects to share information with a party, she can often exercise control over the information flow. Take, for example, Example Social, a popular social network. The user may decide she does not like Example Social's privacy or security practices, so she does not visit examplesocial.com. But if Example Social provides a social sharing widget embedded in another website, the user may be unaware she is giving information to Example Social and unable to exercise control over the information flow.

Second, we recognize that market pressures are an important factor in encouraging good privacy and security practices. If users do not expect that they will share information with an organization, it is unlikely to experience market pressure from users to protect the security and privacy of their information. In practice, moreover, third parties may not experience sufficient market pressure from first parties since increasingly third parties do not have a direct business relationship with the first party websites they appear on. We therefore require a greater degree of user control over information sharing with such organizations.

Last, third parties are often in a position to collect a sizeable proportion of a user's browsing history – information that can be uniquely sensitive and easily associated

with a user's identity. We wish to provide user control over such information flows.

We recognize that, unlike with a bright-line rule, there can be close calls in applying our standard for what constitutes a first party or a third party. But we believe that in practice, such close calls will be rare. The overwhelming majority of content on the web can be classified as first party or third party, with few cases of ambiguity in practice.

We require a confidence at a "high probability" before a party can consider itself a first party. Where there is reasonable ambiguity about whether a user has intentionally interacted with a party, it must consider itself a third party. Our rationale is that, in the rare close cases, a website is in the best position to understand its users' expectations. We therefore impose the burden of understanding user expectations on the website. We also wish, in close cases, to err on the side of conforming to user expectations and protecting user privacy. If the standard is insufficiently protective, ordinary users have limited recourse; if the standard imposes excessive limits, websites retain the safety valve of explicitly asking for user permission.

2.3.2.2 Common Examples and Use Cases

1. A user accesses an Example News article. The page includes an advertisement slot, which loads content from many companies other than Example News. Those companies are third parties.
2. A user accesses an Example News article. The page includes an analytics script that is hosted by Example Analytics, an analytics service. Example Analytics is a third party.
3. A user accesses an Example News article. It includes a social sharing widget from Example Social, a popular social network. Example Social is a third party.
4. A user visits Example Diary, which is hosted by the free blogging service Example Blog Hosting but located at examplediary.com. Example Blog Hosting is a third party.
5. A user launches Example Application, an app on a mobile device. The app includes a library from Example Advertising Network that displays ads. Example Advertising Network is a third party.

2.3.2.3 Multiple First Parties

There will almost always be only one party that the average user would expect to communicate with: the provider of the website the user has visited. But, in rare cases, users may expect that a website is provided by more than one party. For example, suppose Example Sports, a well known sports league, collaborates with Example Streaming, a well known streaming video website, to provide content at www.examplesportsonexamplestreaming.com. The website is prominently advertised and branded as being provided by both Example Sports and Example Streaming. An ordinary user who visits the website may recognize that it is operated by both Example Sports and Example Streaming.

2.3.2.4 User Interaction with Third-Party Content

A party may start out as a third party but become a first party later on, after a user interacts with it. If content from a third party is embedded on a first party page, the third party may become an additional first party if it can infer with high probability that the average user knowingly and intentionally communicated with it. If a user merely moused over, closed, or muted third-party content, the party would not be able to draw such an inference.

2.3.2.4.1 Examples and Use Cases

Example: Example Weather offers an unbranded weather widget that is embedded into websites, including Example News. The widget contains small links to Example Weather's website and privacy policy. A user visits Example News and scrolls through the weekly forecast in the Example Weather widget.

Discussion: Example Weather is a third party. The user has interacted with Example Weather's widget, but an ordinary user would not expect that scrolling through the widget involves communicating with Example News.

Example: Example Social, a popular social network, hosts a social sharing button that other websites can embed. The button is colored and styled in the same fashion as Example Social's website, contains descriptive text that is specific to Example Social, includes Example Social's logo, and very frequently appears on Example Social's website. Example News embeds the Example Social button, and a user clicks it.

Discussion: Example Social is a first party once the user clicks its embedded social sharing button. The average user would understand that by clicking the button she is communicating with Example Social.

3. Information Practices

3.1 Reception, Retention, Use, and Sharing

A party ***receives*** data if the data comes within its control.

A party ***retains*** data if the data remains within the party's control.

A party ***uses*** data if the party processes the data for any purpose, including for storage.

A party ***shares*** data if the party enables another party to receive the data.

3.2 First Party

A first party MUST NOT share information with a third party that the third party is prohibited from receiving itself.

Best Practice 1: Additional Voluntary Measures

Best Practice 1: Additional Voluntary Measures

A first party may voluntarily take steps to protect user privacy when responding to a Do Not Track request.

3.3 Third Party

3.3.1 General Rule

A third party MUST NOT receive, retain, use, or share any information related to communication with a user or user agent. There are exceptions to this general rule as defined in the following sections. In case of ambiguity, an exception MUST be construed narrowly. Each exception operates independently; exceptions cannot be combined except where explicitly noted otherwise.

3.3.2 Exceptions

3.3.2.1 Protocol Information

3.3.2.1.1 Definition

Protocol information includes:

- any information that a user agent necessarily shares with a web server when it communicates with the web server (e.g. IP address and User-Agent), and
- the URL of the top-level page, communicated via a Referer header or other means, unless the URL contains information that is not unlinkable (e.g. a username or user ID).

Protocol information does not include:

- any information that a web server could cause to not be sent but still communicate with the user agent (e.g. a cookie or a Request-URI parameter generated by the user agent), except the URL of the top-level page, and
- any data added by a network intermediary that the operator of a web server has actual knowledge of (e.g. a unique device identifier HTTP header).

3.3.2.1.2 In General

A third party MAY receive and use protocol information for any purpose, subject to a two-week retention period.

3.3.2.1.3 Non-Normative Discussion: Contextual Personalization

Under the general rule on protocol information a third party MAY temporarily use a top-level page URL for the purpose of contextually personalizing content.

3.3.2.1.4 Additional Limit on Geolocation

Under the general rule a third party MAY temporarily use an IP address for geolocation. The geolocation MUST be coarse.

3.3.2.1.5 Security and Fraud Prevention

A third party MAY receive and use protocol information for the detection and prevention of security breaches and fraudulent activity, subject to a six-month retention period and the restrictions imposed in the subsequent sections on security and fraud prevention.

3.3.2.2 Unlinkable Data

3.3.2.2.1 Definitions

A dataset is ***unlinkable*** when there is a high probability that it contains only information which could not be linked to a particular user, user agent, or device by a skilled analyst.

N-unlinkability is the special case of K-anonymity where all values are considered part of the pseudo-identifier.

3.3.2.2.2 Validation

Third parties that receive, retain, or use unlinkable data MUST either:

1. publicly publish information that is sufficiently detailed for a skilled analyst to evaluate the implementation, or
2. ensure that any datasets are at least 1024-unlinkable.

3.3.2.2.3 Information That Is Unlinkable When Received

A third party MAY receive non-protocol information if it is, independent of protocol information, unlinkable data. The data MAY be retained and used subject to the same

limitations as protocol information. Such data MUST be disassociated from protocol information when it is first used or within two weeks, whichever is sooner.

Example: Example Advertising sets a language preference cookie that takes on few values and is shared by many users. Log entries containing this preference cookie as well as protocol information are collected on each of Example Advertising's webservers. When Example Advertising processes its logs, it computes unlinkable datasets using the protocol logs and language cookies. After that process, Example Advertising no longer stores log files that associate protocol log entries with the language cookies.

3.3.2.2.4 Information That Is Unlinkable After Aggregation

During the period in which a third party may use protocol information for any purpose, it may aggregate protocol information and unlinkable data into an unlinkable dataset. Such a dataset may be retained indefinitely and used for any purpose.

Example: Example Advertising maintains a dataset of how many times per week Italy-based users load each of its ads on Example News.

3.3.2.3 Outsourcing

A first party MAY outsource website functionality to a third party, in which case the third party may act as the first party under this standard with the following additional restrictions.

3.3.2.3.1 Technical Precautions

3.3.2.3.1.1 Operative Text

Throughout all data reception, retention, and use, outsourced service providers MUST use all feasible technical precautions to both mitigate the linkability of and prevent the linking of data from different first parties.

Structural separation ("siloing") of data per first party, including both

1. separate data structures and
2. avoidance of shared unique identifiers

are necessary, but not necessarily sufficient, technical precautions.

3.3.2.3.1.2 Non-Normative Discussion

3.3.2.3.1.2.1 Siloing in the Browser

Outsourcing services should use browser access control features so that stored data specific to one first party is never accessed or received when the user visits another first party.

3.3.2.3.1.2.1.1 Same-Origin Policy

The same-origin policy silos stored data by domain name. An outsourcing service can use a different domain name for each first party.

Example: Example Analytics provides an outsourced analytics service to Example News and Example Sports, two unrelated websites. Example Analytics stores its cookies for Example News at examplenews.exampleanalytics.com, and it stores its cookies for Example Sports at examplesports.exampleanalytics.com.

3.3.2.3.1.2.1.2 Cookie Path Attribute

The HTTP cookie path can be used to silo data to a first party.

Example: Example Analytics stores its cookies for Example News with "Path=/examplenews", and it stores its cookies for Example Sports with "Path=/examplesports".

3.3.2.3.1.2.1.3 Storage Key

For key/value storage APIs, such as Web Storage and Indexed Database, an outsourcing service can use a different key or key prefix for each first party.

Example: Example Analytics stores data for Example News at window.localStorage["examplenews"] and data for Example Sports at window.localStorage["examplesports"].

3.3.2.3.1.2.2 Siloing in the Backend

3.3.2.3.1.2.2.1 Encryption Keys

An outsourcing service should encrypt each first party's data with a different set of keys.

3.3.2.3.1.2.2.2 Access Controls

An outsourcing service should deploy access controls so that only authorized personnel are able to access siloed data, and only for authorized purposes.

3.3.2.3.1.2.2.3 Access Monitoring

An outsourcing service should deploy access monitoring mechanisms to detect improper use of siloed data.

3.3.2.3.1.2.3 Retention in the Backend

An outsourcing service should retain information only so long as necessary to provide necessary functionality to a first party. If a service creates periodic reports, for example, it should delete the data used for a report once it is generated. An outsourcing service should be particularly sensitive to retaining protocol logs, since they may allow correlating user activity across multiple first parties.

3.3.2.3.2 Internal Practices

3.3.2.3.2.1 Operative Text

Throughout all data reception, retention, and use, outsourced service providers MUST use sufficient internal practices to prevent the linking of data from different first parties.

3.3.2.3.2.2 Non-Normative Discussion

3.3.2.3.2.2.1 Policy

An outsourcing service should establish a clear internal policy that gives guidance on how to receive, retain, and use outsourced data in compliance with this standard.

3.3.2.3.2.2.2 Training

Personnel that interact with outsourced data should be familiarized with internal policy on compliance with this standard.

3.3.2.3.2.2.3 Supervision and Reporting

An outsourcing service should establish a supervision and reporting structure for detecting improper access.

3.3.2.3.2.2.4 Auditing

External auditors should periodically examine an outsourcing service to assess whether it is in compliance with this standard and has adopted best practices. Auditor reports should be made available to the public.

3.3.2.3.3 Use Direction

An outsourced service

1. MUST use data retained on behalf of a first party ONLY on behalf of that first party, and
2. MUST NOT use data retained on behalf of a first party for their own business purposes, or for any other reasons.

3.3.2.3.4 First-Party Requirements

3.3.2.3.4.1 Representation

A first party's representation that it is in compliance with this standard includes a representation that its outsourcing service providers comply with this standard.

3.3.2.3.4.2 Contract

A first party MUST enter into a contract with an outsourcing service provider that requires that outsourcing service provider to comply with these requirements.

3.3.2.4 User Permission

A website my engage in practices otherwise prohibited by this standard if a user grants permission. Permission may be attained through the browser API defined in the companion Tracking Preference Expression document. A website may also rely on "out-of-band" consent attained through a different technology. An "out-of-band" choice mechanism has the same effect under this standard as the browser exception API,

provided that it satisfies the following bright-line requirements:

1. **Actual presentation:** The choice mechanism MUST be actually presented to the user. It MUST NOT be on a linked page, such as a terms of service or privacy policy.
2. **Clear terms:** The choice mechanism MUST use clear, non-confusing terminology.
3. **Independent choice:** The choice mechanism MUST be presented independent of other choices. It MUST NOT be bundled with other user preferences.
4. **No default permission:** The choice mechanism MUST NOT have the user permission preference selected by default.

An "out-of-band" choice mechanism must additionally satisfy the following high-level standard:

An ordinary user would know that the choice overrides his or her privacy protections under this standard.

3.3.2.5 Security

3.3.2.5.1 Operative Text

A third party MAY receive, retain, and use data about a particular user or user agent for the purpose of ensuring its security, provided that there are reasonable grounds to believe the user or user agent was attempting to breach the party's security at the time the data was received.

Note: This draft does not address the extent to which technical and business precautions are required for security data.

3.3.2.5.2 Non-Normative Discussion

This exception grants third parties (e.g. advertising networks) some latitude to mitigate security risks. Websites that users store sensitive personal information on (e.g. financial services and webmail) are all first-party; they are able to receive, retain, and use information about all users for security purposes.

3.3.2.6 Fraud Prevention

3.3.2.6.1 Operative Text

A third party MAY receive, retain, and use data about a particular user or user agent for the purpose of preventing fraud, provided that there are reasonable grounds to believe the user or user agent was attempting to commit fraud at the time the data was received.

Note: This draft does not address the extent to which technical and business precautions are required for fraud prevention data.

3.3.2.6.2 Non-Normative Discussion

When a user meaningfully interacts with third-party content (e.g. clicking an ad), the third party can receive, retain, and use data for fraud prevention. Third parties can also use protocol information for fraud prevention. This exception provides an additional capability to, in certain circumstances, track impressions for fraud prevention.

3.3.2.7 Unknowing Information Practices

Note: This section was recently added and has not been extensively discussed with stakeholders. Please consider it a preliminary position.

A party MAY receive, retain, and use data as otherwise prohibited by this standard, so long as is unaware of such information practices and has made reasonable efforts to understand its information practices. If a party learns that it possesses information in violation of this standard, it MUST delete that information at the earliest practical opportunity.

A. References

A.1 Normative references

No normative references.

A.2 Informative references

No informative references.

Do Not Track — Combined Proposal

Unofficial Draft 19 June 2012

Editor:
Aleecia M. McDonald, Mozilla

Abstract

This partial draft attempts to reflect some of the points of consensus across multiple proposals presented at the Washington, DC f2f meeting, including issue-10 (what is a first party), issue-17 (data use by a first party), issue-19 (data collection and use by a third party), issue-31 (data minimization), issue-49 (third party on behalf of first party), and issue-73 (silo v. contract for outsourcing partner). It does not get into the disputed areas of issue-22 (operational use), issue-24 (fraud exemption), issue-25 (research exemption).

If you remember the eight page poster on the wall with only two points of major disagreement in DC, this draft tries to capture the parts where we basically agreed. The idea is to have a format we can now go through and tweak and fix, then hand off to the Compliance editors to integrate into a public draft. Our next step is to see where we have consensus on these issues, where we are still working out wording, and where there are differences. Please review this text carefully prior to discussions during the Seattle f2f.

Much of the text comes from two proposals, one from Jonathan / Peter / Tom, the other from Shane et. al. Extraordinarily little text is actually new. The structure is different, with definitions pulled to the beginning and a simplified hierarchy.

Status of This Document

This document is merely a public working draft of a potential specification. It has no official standing of any kind and does not represent the support or consensus of any standards organisation.

Table of Contents

1. Definitions

A **functional entity** is any commercial, nonprofit, or governmental organization, a subsidiary or unit of such an organization, or a person.

Functional entities are **affiliated** when they are related by both common majority ownership and common control.

A **party** is a set of functional entities that are affiliated and follow requirements to be easily discoverable.

A **network interaction** is an HTTP request and response, or any other set of logically related network traffic.

An **outsourced party** is any party, in a specific network interaction, that is working on behalf of a specific first or third party in compliance with the outsourced party information practices.

A **first party** is any party in a specific network interaction that can infer with high probability that the user knowingly and intentionally communicated with it, and complies with DNT under first party information practices.

A **third party** is any party in a specific network interaction that is not a first party, an outsourced party, or the user. A party without the ability to infer with high probability that the user knowingly and intentionally communicated with it MUST represent itself as a third party and comply with third party information practices. A first party MAY choose to represent itself as a third party and comply with third party information practices.

A dataset is **unidentifiable** when there is a high probability that it contains only information which could not be linked to a particular user, user agent, or device by a skilled analyst. N-unlinkability is the special case of K-anonymity where all values are considered part of the pseudo-identifier.

Protocol information includes:

- any information that a user agent necessarily shares with a web server when it communicates with the web server (e.g. IP address and User-Agent), and
- the URL of the top-level page, whether communicated via a Referer header or other means.

Protocol information does not include:

- any information that a web server could cause to not be sent but still communicate with the user agent (e.g. a cookie or a Request-URI parameter generated by the user agent), except the URL of the top-level page, and
- any data added by a network intermediary that the operator of a web server has actual knowledge of (e.g. a unique device identifier HTTP header).

A party **collects** data if the data comes within its control and the control of that data is not transient.

Note

Open action on defining collection; this is not done

A party **retains** data if the data remains within the party's control.

A party **uses** data if the party processes the data for any purpose, including for storage.

A party **shares** data if the party enables another party to collect the data.

---

2. Information Practices for All Parties

This section of the specification applies to all parties who comply with an incoming DNT signal. See the companion [*TRACKING-DNT*] document for information on how to respond to an incoming signal.

2.1 Additional Voluntary Measures

This specification sets a minimum common standard for compliance. Any party MAY take additional steps to protect user privacy when responding to a Do Not Track request.

2.2 User Permission and Consent

Note

We have discussed this at length in action-152, action-157, and action-159. This section attempts to merge several texts together.

Sites MAY override a user's DNT preference if they have received explicit, informed consent to do so.

A party may engage in practices otherwise prohibited by this standard if a user grants permission. When seeking an exemption, sites SHOULD communicate those requests clearly, accurately, and in line with consumer protection laws in the jurisdictions in which they operate. Permission may be attained through either (A) the browser API defined in the companion [*TRACKING-DNT*] document or (B) an "out-of-band" consent attained through a different technology. An "out-of-band" choice mechanism has the same effect under this standard as the browser exception API.

A party may receive multiple conflicting signals from users. Specific permission overrides a general permission. If a party has received prior consent for tracking a given user, user agent, or device, that consent overrides the general preference indicated by the DNT header field. If a party chooses to track based on that prior consent, the corresponding tracking status MUST indicate that tracking is occurring based on consent and SHOULD supply a link to a means for the user to remove or modify that consent, as described in [*TRACKING-DNT*]. For example, if you have two conflicting signals from a user with a global DNT:1 yet also have the user's consent specific to you (via browser API or out-of-band consent) then the user's consent is the correct signal to follow.

An "out-of-band" choice mechanism MUST additionally satisfy the following condition: An ordinary user would know that the choice overrides his or her general Tracking Preference.

2.2.1 Non-normative

This section is non-normative.

Many organizations have developed direct consent mechanisms for web-wide tracking prior to this standard. Interactions with users to obtain consent may be contextual. For example, if a service has an obvious cross-site tracking function that the user deliberately signs up for then this could be deemed to have achieved “explicit and informed” consent from a user without directly addressing its relation to a Tracking Preference (which wasn’t contemplated at the time the consent experience was designed). Even in these cases, organizations should recall that users with DNT:1 are requesting privacy, and strongly consider providing Tracking Preference references in associated product or service materials such as a privacy policy, help center, or a separate notice to users who have indicated preferences for privacy. Organizations should also scope the user's consent to the task the user is likely to understand. For example, a user who signed up for a news tracking service that would note which articles he or she read across the entire web in order to suggest relevant articles in the future might consider just using the service as sufficient contextual consent for gathering information about news. However, that does not mean the user has also consented to that information being used in any other context or by other parties. Contextual consent is specific to a product or feature, as understood by the user.

Companies should not seek to obtain explicit, informed consent from users in non-obvious ways such as placing these details in their Terms of Service or their Privacy Center if it will not be obvious to users that the nature of the service will lead the company to ignore a user’s Tracking Preference based on the nature of the consent the user is granting. As an example where context would be insufficient to establish consent, a company could not obtain explicit, informed consent to ignore DNT in third-party settings by placing notice in a privacy policy or the terms of use for an email service or a game if that company also happens to own an advertising network. The company would need to present users with a request for consent. It is a good practice to present a choice that signing up for a service will override a Tracking Preference setting, or the option to decline the service, at the time the user's choice is being made. The W3C geolocation API is one example of out-of-band consent, which we discuss in the section on geolocation.

Out-of-band consent will be further reinforced in user interactions through either the Header Response or Well-Known URI approaches to replying to user Tracking Preferences. This will provide a constant reminder of prior consent on each interaction and provide a resource (link) to allow the user to understand how this consent was achieved and ideally present options to alter that consent if the user chooses to do so.

We suggest the following properties for any out-of-band consent mechanism:

Actual presentation

The choice mechanism must be actually presented to the user. It must not only be on a linked page, such as a terms of service or privacy policy.

Clear terms

The choice mechanism must use clear, non-confusing terminology.

Independent choice

The choice mechanism must be presented independent of other choices. It must

not be bundled with other user preferences.

No default permission
: The choice mechanism must not have the user permission preference selected by default.

What constitutes explicit consent is not necessarily the same across all legal jurisdictions. It is the site's responsibility to ensure that it has consent.

2.3 Unidentifiable Data

Any party MAY collect, retain, or use unidentifiable data, subject to the requirements that the party MUST either:

1. publicly publish information that is sufficiently detailed for a skilled analyst to evaluate the implementation, or
2. ensure that any datasets are at least 1024-unlinkable.

Unidentifiable information will either be unidentifiable at the time of collection, or be made unidentifiable by aggregating data after it is collected; both are described below.

2.3.1 Information That Is Unidentifiable When Collected

A party may collect non-protocol information if it is, independent of protocol information, unidentifiable data. The data may be retained and used subject to the same limitations as protocol information.

Example

```
Example Advertising sets a language preference cookie that stores one of a
few values, such as us, de, and so on, and thousands of users share the
same language preference.
```

2.3.2 Information That Is Unidentifiable After Aggregation

Note

If we do not adopt the notion of a grace period for log files, then this section applies only to parties that know they are always first parties. For everyone else, the only form of permitted aggregation will be at the point of collection as things are currently written.

During the period in which a party may use protocol information prior to processing, it may aggregate protocol information and unidentifiable data into an unidentifiable dataset. Such a dataset may be retained indefinitely and used for any purpose.

Example

```
Example Advertising maintains a dataset of how many times per week
Italy-based users load an ad on Example News.
```

2.3.3 Non-normative

This section is non-normative.

Note

It would be helpful to describe what 1024-unlinkable means so we do not send implementers scrambling for other texts. Also useful, a discussion that this can be either calculated mathematically (with a pointer to a readable reference) or by estimating based on actual data, perhaps with non-DNT users or pre-DNT users.

2.4 Discoverability

Note

While there is disagreement over whether discoverability is sufficient, we do seem to be converging on what discoverability is. Should we decide discoverability is not helpful, or insufficient, we can easily remove this section.

A functional entity must make its affiliated functional entities easily discoverable by a user.

2.4.1 Non-Normative Discussion

This section is non-normative.

Affiliation may be made easily discoverable by a user in many ways, including but not limited to: prominent and common branding on pages, one click away within a privacy policy, or a machine-readable format in a well-known location. As a general guideline: if a lawsuit could be brought against two different entities, they are not the same functional entity. Similarly, if two portions of a legal entity have different privacy policies, they should not be considered the same functional entity under Do Not Track.

2.5 Additional Requirements Based on Party Status

In addition to the information practices for all parties as described in this section, for each network interaction an additional set of information practices applies based on which type of party you are during that network interaction: first party, outsourced party, or third party.

2.5.1 Non-Normative Discussion

This section is non-normative.

Determination of a party's status is limited to a single transaction because a party's status may be affected by time, context, or any other factor that influences user expectations.

Other than some third parties becoming first parties when users interact (for example, social widgets or ads,) party status will usually stay stable for the entire interaction with a given user.

---

3. Information Practices for First Parties

This section of the document applies just to first parties.

A first party MUST NOT share information with a third party that the third party is prohibited from collecting itself. A first party MUST NOT share (send or receive) identifiable information about a user to any party it does not have an outsource relationship with.

Note

While confining data just to the first party reflects discussions of the TPWG to date, newer members have concerns about these provisions.

3.1 Non-Normative Discussion

This section is non-normative.

3.1.1 Overview

We draw a distinction between those parties an ordinary user would or would not expect to share information with, "first parties" and "third parties" respectively. The delineation exists for three reasons.

First, when a user expects to share information with a party, she can often exercise control over the information flow. Take, for example, Example Social, a popular social network. The user may decide she does not like Example Social's privacy or security practices, so she does not visit examplesocial.com. But if Example Social provides a social sharing widget embedded in another website, the user may be unaware she is giving information to Example Social and unable to exercise control over the information flow.

Second, we recognize that market pressures are an important factor in encouraging good privacy and security practices. If users do not expect that they will share information with an organization, it is unlikely to experience market pressure from users to protect the security and privacy of their information. In practice, moreover, third parties may not experience sufficient market pressure from first parties since increasingly third parties do not have a direct business relationship with the first party websites they appear on. We therefore require a greater degree of user control over information sharing with such organizations.

Last, third parties are often in a position to collect a sizeable proportion of a user's browsing history — information that can be uniquely sensitive and easily associated with a user's identity. We wish to provide user control over such information flows.

We recognize that, unlike with a bright-line rule, there can be close calls in applying our

standard for what constitutes a first party or a third party. But we believe that in practice, such close calls will be rare. The overwhelming majority of content on the web can be classified as first party or third party, with few cases of ambiguity in practice.

We require a confidence at a "high probability" before a party can consider itself a first party. Where there is reasonable ambiguity about whether a user has intentionally interacted with a party, it must consider itself a third party. Our rationale is that, in the rare close cases, a website is in the best position to understand its users' expectations. We therefore impose the burden of understanding user expectations on the website. We also wish, in close cases, to err on the side of conforming to user expectations and protecting user privacy. If the standard is insufficiently protective, ordinary users have limited recourse; if the standard imposes excessive limits, websites retain the safety valve of explicitly asking for user permission.

3.1.2 Common Examples and Use Cases

1. A user accesses an Example News article. The page includes an advertisement slot, which loads content from many companies other than Example News. Those companies are third parties.
2. A user accesses an Example News article. The page includes an analytics script that is hosted by Example Analytics, an analytics service. Example Analytics is a third party.
3. A user accesses an Example News article. It includes a social sharing widget from Example Social, a popular social network. Example Social is a third party.
4. A user visits Example Diary, which is hosted by the free blogging service Example Blog Hosting but located at examplediary.com. Example Blog Hosting is a third party.
5. A user launches Example Application, an app on a mobile device. The app includes a library from Example Advertising Network that displays ads. Example Advertising Network is a third party.

3.1.3 Multiple First Parties

There will almost always be only one party that the average user would expect to communicate with: the provider of the website the user has visited. But, in rare cases, users may expect that a website is provided by more than one party. For example, suppose Example Sports, a well known sports league, collaborates with Example Streaming, a well known streaming video website, to provide content at www.examplesportsonexamplestreaming.com. The website is prominently advertised and branded as being provided by both Example Sports and Example Streaming. An ordinary user who visits the website may recognize that it is operated by both Example Sports and Example Streaming.

---

4. Information Practices for Third Parties

This section of the document applies just to third parties.

4.1 General Rule

A third party may not collect, retain, use, or share any information related to communication with a user or user agent. There are exceptions to this general rule as defined in the following sections.

4.2 Permitted Uses

We recognize a limited set of data uses as important enough to continue even with data in potentially identifiable form. For all other uses, many of which are quite valuable, sites can ask users for permission. Note that first parties are not constrained to permitted uses. Outsourced parties may act as the party they work with would act: if working with a third party, they are also bound to this list of permitted uses.

Note

While we agree on this general structure, we do not agree on uses at this point.

4.3 User Interaction with Third Party Content

A party may start out as a third party but become a first party later on, after a user interacts with it. If content from a third party is embedded on a first party page, the third party may become an additional first party if it can infer with high probability that the average user knowingly and intentionally communicated with it. If a user merely moused over, closed, or muted third party content, the party would not be able to draw such an inference.

4.3.1 Examples and Use Cases

This section is non-normative.

Example: Example Weather offers an unbranded weather widget that is embedded into websites, including Example News. The widget contains small links to Example Weather's website and privacy policy. A user visits Example News and scrolls through the weekly forecast in the Example Weather widget.

Discussion: Example Weather is a third party. The user has interacted with Example Weather's widget, but an ordinary user would not expect that scrolling through the widget involves communicating with Example News.

Example: Example Social, a popular social network, hosts a social sharing button that other websites can embed. The button is colored and styled in the same fashion as Example Social's website, contains descriptive text that is specific to Example Social, includes Example Social's logo, and very frequently appears on Example Social's website. Example News embeds the Example Social button, and a user clicks it.

Discussion: Example Social is a first party once the user clicks its embedded social sharing button. The average user would understand that by clicking the button she is communicating with Example Social.

---

5 Information Practices for Outsourcing

5. Information Practices for Outsourcing

This section applies to parties engaging in an outsourcing relationship, wherein one party "stands in the shoes" of another party to perform a specific task. Both parties have responsibilities, as detailed below.

A first party or a third party MAY outsource functionality to another party, in which case the third party may act as the original first party or third party under this standard, with the following additional restrictions:

- Data collected by each outsourced company is separated for each party they collect data for by both technical means and organizational process, AND
- The outsourced company has no independent rights to the collected information, AND
- A contractual relationship exists between the outsourced and the party they collect data for that outlines and mandates these requirements.

An outsourced company acting on the behalf of another party is subject to all of the same restrictions on that party (for First or Third party, as appropriate.)

5.1 Non-Normative

This section is non-normative.

Outsourced companies that act purely as vendors for their customers (often first parties in this context) are not the intended target for the Tracking Preference Expression but it is important there are no unintended activities that are extended to another party through this allowance. In all cases, its expected an outsourced company acting on the part of a customer follows all of the same restrictions placed on that customer.

For the data separation requirement, outsourced companies have technical options to achieve appropriate separation but in each the critical element is that data is never reconstituted for users that have indicated a preference not to be tracked. One possible approach would be to leverage a per partner hash against a common cookie identifier, ensuring the resulting identifier is consistent for a specific customer, but is unable to be linked with another customer's identifier.

Contractual requirements that enforce data rights and responsibilities for separation are a critical element of establishing an outsourcer acting on another party's behalf. Contracts may occur directly through parties (for example, a Publisher in an Ad Network) or between intermediaries (for example, an Ad Network acting through an Ad Exchange). In either case, data separation and removal of independent rights are necessary elements that must survive intermediary contractual constructs.

5.2 Technical Precautions

Throughout all data collection, retention, and use, outsourced parties MUST use all feasible technical precautions to both mitigate the identifiability of and prevent the identification of data from different first parties.

Structural separation ("siloing") of data per first party, including both

1. separate data structures and
2. avoidance of shared unique identifiers

are necessary, but not necessarily sufficient, technical precautions.

5.3 Non-Normative Discussion

This section is non-normative.

5.3.1 Siloing in the Browser

Outsourcing services should use browser access control features so that stored data specific to one party is never accessed or collected when the user visits another party.

5.3.1.1 Same-Origin Policy

The same-origin policy silos stored data by domain name. An outsourcing service can use a different domain name for each first party.

Example

```
Example Analytics provides an outsourced analytics service to Example News
and Example Sports, two unrelated websites. Example Analytics stores its
cookies for Example News at examplenews.exampleanalytics.com, and it
stores its cookies for Example Sports at
examplesports.exampleanalytics.com.
```

5.3.1.2 Cookie Path Attribute

The HTTP cookie path can be used to silo data to a first party.

Example

```
Example Analytics stores its cookies for Example News with
"Path=/examplenews", and it stores its cookies for Example Sports with
"Path=/examplesports".
```

5.3.1.3 Storage Key

For key/value storage APIs, such as Web Storage and Indexed Database, an outsourcing service can use a different key or key prefix for each first party.

Example

```
Example Analytics stores data for Example News at
window.localStorage["examplenews"] and data for Example Sports at
window.localStorage["examplesports"].
```

5.3.2 Siloing in the Backend

5.3.2.1 Encryption Keys

An outsourcing service should encrypt each first party's data with a different set of keys.

5.3.2.2 Access Controls

An outsourcing service should deploy access controls so that only authorized personnel are able to access siloed data, and only for authorized purposes.

5.3.2.3 Access Monitoring

An outsourcing service should deploy access monitoring mechanisms to detect improper use of siloed data.

5.3.3 Retention in the Backend

An outsourcing service should retain information only so long as necessary to provide necessary functionality to a first party. If a service creates periodic reports, for example, it should delete the data used for a report once it is generated. An outsourcing service should be particularly sensitive to retaining protocol logs, since they may allow correlating user activity across multiple first parties.

5.4 Internal Practices

Throughout all data collection, retention, and use, outsourced parties MUST use sufficient internal practices to prevent the identification of data from different parties.

5.4.1 Non-Normative Discussion

This section is non-normative.

5.4.1.1 Policy

An outsourcing service should establish a clear internal policy that gives guidance on how to collect, retain, and use outsourced data in compliance with this standard.

5.4.1.2 Training

Personnel that interact with outsourced data should be familiarized with internal policy on compliance with this standard.

5.4.1.3 Supervision and Reporting

An outsourcing service should establish a supervision and reporting structure for detecting improper access.

5.4.1.4 Auditing

External auditors should periodically examine an outsourcing service to assess whether it is in compliance with this standard and has adopted best practices. Auditor reports should be made available to the public.

5.5 Use Direction

An outsourced service:

1. MUST use data retained on behalf of a party ONLY on behalf of that party, and
2. MUST NOT use data retained on behalf of a party for their own business purposes, or for any other reasons.

5.6 First Party or Third Party Requirements

5.6.1 Representation

A party's representation that it is in compliance with this standard includes a representation that its outsourcing parties comply with this standard.

5.6.2 Contract

A first party MUST enter into a contract with an outsourced party that requires that outsourced party to comply with these requirements.

A. References

A.1 Normative references

[TRACKING-DNT]
Roy T. Fielding; David Singer. *Tracking Preference Expression (DNT).* 13 March 2012. W3C Working Draft. (Work in progress.) URL: http://www.w3.org/TR/2011/WD-tracking-dnt-20120313/

A.2 Informative references

No informative references.

Tracking Compliance and Scope Specification

W3C Editor's Draft 23 May 2012

This version:
http://www.w3.org/2011/tracking-protection/drafts/tracking-compliance.html
Latest published version:
http://www.w3.org/TR/tracking-compliance/
Latest editor's draft:
http://www.w3.org/2011/tracking-protection/drafts/tracking-compliance.html
Previous version:
http://www.w3.org/TR/2012/WD-tracking-compliance-20120313/
Editors:
Justin Jeffrey Brookman, CDT
Sean Harvey, Google
Erica Newland, CDT
Heather West, Google

Abstract

This specification defines the meaning of a Do Not Track (DNT) preference and sets out practices for websites to comply with this preference.

Status of This Document

This section describes the status of this document at the time of its publication. Other documents may supersede this document. A list of current W3C publications and the latest revision of this technical report can be found in the W3C technical reports index at http://www.w3.org/TR/.

Note

This document is a snapshot of live discussions within the Tracking Protection Working Group. It does not yet capture all of our work. For example, we have issues that are [PENDING REVIEW] with complete text proposals that did not make it into this draft. Text in white is typically [CLOSED]: we have reached a consensus decision. Text in blue boxes presents multiple options the group is considering. In

some cases we are close to agreement, and in others we have more to discuss. An issue tracking system is available for recording raised, open, pending review, closed, and postponed issues regarding this document.

Note

We have not yet reviewed comments from the Community Group associated with this work. We thank them for their time and detailed feedback, and will address their comments in the near future.

This document was published by the Tracking Protection Working Group as an Editor's Draft. If you wish to make comments regarding this document, please send them to public-tracking@w3.org (subscribe, archives). All feedback is welcome.

Publication as an Editor's Draft does not imply endorsement by the W3C Membership. This is a draft document and may be updated, replaced or obsoleted by other documents at any time. It is inappropriate to cite this document as other than work in progress.

This document was produced by a group operating under the 5 February 2004 W3C Patent Policy. W3C maintains a public list of any patent disclosures made in connection with the deliverables of the group; that page also includes instructions for disclosing a patent. An individual who has actual knowledge of a patent which the individual believes contains Essential Claim(s) must disclose the information in accordance with section 6 of the W3C Patent Policy.

Table of Contents

1. Introduction

The World Wide Web (WWW, or Web) consists of millions of sites interconnected through the use of hypertext. Hypertext provides a simple, page-oriented view of a wide variety of information that can be traversed by selecting links, manipulating controls, and supplying data via forms and search dialogs. A Web page is usually composed of many different information sources beyond the initial resource request, including embedded references to stylesheets, inline images, javascript, and other elements that might be automatically requested as part of the rendering or behavioral processing defined for that page.

Each of the hypertext actions and each of the embedded resource references might refer to any site on the Web, leading to a seamless interaction with the user even though the pages might be composed of information requested from many different and possibly independent Web sites. From the user's perspective, they are simply visiting and interacting with a single brand — the **first-party** Web property — and all of the technical details and protocol mechanisms that are used to compose a page representing that brand are hidden behind the scenes.

It has become common for Web site owners to collect data regarding the usage of their sites for a variety of purposes, including what led the user to visit their site (referrals), how effective the user experience is within the site (web analytics), and the nature of who is using their site (audience segmentation). In some cases, the data collected is used to dynamically adapt the content (personalization) or the advertising presented to the user (targeted advertising). Data collection can occur both at the first-party site and via third-party providers through the insertion of tracking elements on each page. A survey of these techniques and their privacy implications can be found in [*KnowPrivacy*].

People have the right to know how data about them will be collected and how it will be used. Empowered with that knowledge, individuals can decide whether to allow their online activities to be tracked and data about them to be collected. Many Internet companies use data gathered about people's online activities to personalize content and target advertising based on their perceived interests. While some people appreciate this personalization of content and ads in certain contexts, others are troubled by what they perceive as an invasion of their privacy. For them, the benefit of personalization is not worth their concerns about allowing entities with whom they have no direct relationship to amass detailed profiles about their activities.

Therefore, users need a mechanism to express their own preference regarding tracking that is both simple to configure and efficient when implemented. In turn, Web sites that are unwilling or unable to offer content without such targeted advertising or data collection need a mechanism to indicate those requirements to the user and allow them

(or their user agent) to make an individual choice regarding user-granted exceptions.

This specification defines the terminology of tracking preferences, the scope of its applicability, and the requirements on compliant first-party and third-party participants when an indication of tracking preference is received. This specification defines the meaning of a Do Not Track preference and sets out practices for websites and other online companies to comply with this preference.

A companion document, [[!!TRACKING-DNT]], defines the HTTP request header field DNT for expressing a tracking preference on the Web, a well-known location (URI) for providing a machine-readable tracking status resource that describes a service's DNT compliance, the HTTP response header field Tk for resources to communicate their compliance or non-compliance with the user's expressed preference, and JavaScript APIs for determining DNT status and requesting a site-specific, user-granted exception.

Issue

ISSUE-117: Terms: tracking v. cross-site tracking
The WG has not come to consensus regarding the definition of tracking and whether the scope of DNT includes all forms of user-identifying data collection or just cross-site data collection/use. This issue will be resolved in the TCS document, though its resolution is a necessary prerequisite to understanding and correctly implementing the protocol defined by this document.

2. Scope and Goals

2.1 Goals

Note

This section consists of proposed text that is meant to address ISSUE-6 and is in active discussion. Currently, it satisfies no one. We will revisit and finalize once the document is more complete.

Issue

ISSUE-6: What are the underlying concerns? Why are we doing this?

While there are a variety of business models to monetize content on the web, many rely on advertising. Advertisements can be targeted to a particular user's interests based on information gathered about one's online activity. While the Internet industry believes many users appreciate such targeted advertising, as well as other personalized content, there is also an understanding that some people find the practice intrusive. If this opinion becomes widespread, it could undermine the trust necessary to conduct business on the Internet. This Compliance specification and a companion [*TRACKING-DNT*] specification are intended to give users a means to indicate their tracking preference and to spell out the obligations of compliant websites that receive the Do Not Track message. The goal is to provide the user with choice, while allowing practices necessary for a smoothly functioning Internet. This should be a win-win for business and consumers alike. The Internet brings millions of users and web sites together in a vibrant and rich ecosystem.

As the sophistication of the Internet has grown, so too has its complexity which leaves all but the most technically savvy unable to deeply understand how web sites collect and use data about their online interactions. While on the surface many web sites may appear to be served by a single entity, in fact, many web sites are an assembly of multiple parties coming together to power a user's online experience. As an additional privacy tool, this specification provides both the technical and compliance guidelines to enable the online ecosystem to further empower users with the ability to communicate a tracking preferences to a web site and its partners.

The accompanying [*TRACKING-DNT*] recommendation explains how a user, through a user agent, can clearly express a desire not to be tracked. This Tracking Compliance and Scope recommendation sets the standard for the obligations of a website that receives such a DNT message.

Taken together these two standards should have four substantial outcomes:

1. Empower users to manage their preference around the collection and correlation of data about Internet activities that occur on different sites and spell out the obligations of sites in honoring those preferences when DNT is enabled.
2. Provide an exceedingly straightforward way for users to gain transparency and control over data usage and the personalization of content and advertising on the web.
3. Enable a vibrant Internet to continue to flourish economically by supporting innovative business models while protecting users' privacy.
4. Establish compliance metrics for operators of online services

This solution is intended to be persistent, technology neutral, and reversible by the user. It aims to preserve a vibrant online ecosystem, privacy-preserving secondary data uses necessary to ecommerce, and adequate security measures. We seek a solution that is persistent, technology neutral, and [something that speaks with the ability to opt back in], but that preserves a vibrant online ecosystem, privacy-preserving secondary data uses, and adequate security measures.

3. Definitions

The Definitions section of this document is in a high state of flux as of this draft and contains a large number of controversial text proposals and open issues for the working group. None of these definitions should be considered definitive or final.

Note

The term **permitted use** is used to indicate a restricted set of conditions under which tracking is allowed in spite of the user's DNT preference. The term **user-granted exception** is used when the user has permitted tracking, usually in the form of a site-specific exception, for a given third-party. In general: permitted uses are additional permissions granted by the standard; user-granted exceptions are additional permissions granted by the user. These words are often confused when drafting new text.

3.1 Drafting Notes

Note

ISSUE-97 : A special rule for URL-shortening services remains an open issue and is not addressed in the proposal put forward in 3.2 through 3.4.

Issue

ISSUE-26 : Providing data to 3rd-party widgets &emdash; does that imply consent?

Note

The proposal put forward here does not provided a special rule for widgets. The same first party vs. third party test for static content applies.

3.2 Parties

3.2.1 Option 1: User Expectations

3.2.1.1 Definition

A "party" is any commercial, nonprofit, or governmental organization, a subsidiary or unit of such an organization, or a person, that an ordinary user would perceive to be a discrete entity for purposes of information collection and sharing. Domain names, branding, and corporate ownership may contribute to, but are not necessarily determinative of, user perceptions of whether two parties are distinct.

3.2.1.2 Discussion

This section is non-normative.

Our definition of what constitutes a "party" is guided by ordinary user expectations. We decline to adopt a conflicting approach that would draw a line at domain names, corporate affiliation, or branding, as discussed below.

3.2.1.2.1 Domain Names

In an uncomplicated world, a party might be delineated by domain boundaries. In practice, however, the domain approach can emphasize differences that would not matter to ordinary users and would be restrictive for many business uses. Suppose Example Company hosts dynamic content on example.com and static images on example-static.com. An average user would understand both domains are operated by Example Company, but a domain name distinction would say the two domains are different parties. Using domain names to differentiate parties would also impose an unnecessary choice on large websites of either hosting all their content on a single domain or having some of their content considered third party. By adopting a user expectations standard, we allow a

single website to span multiple domains.

A domain name approach can also gloss over relevant differences from a user expectations perspective. Suppose Example Company hires an analytics company and aliases the domain analytics.example.com to the analytics company's website. By user expectations, and corporate affiliation and branding, the analytics company would be a separate party. Moreover, circumventing the limits imposed by this standard would require nothing more than switching domain names. The user expectations standard we adopt recognizes that multiple parties may exist at a single domain.

3.2.1.2.2 Corporate Affiliation

Corporate families can consist of businesses in completely unrelated industries; users may have limited understanding of how businesses are related by corporate ownership or control. Moreover, by creating affiliates for the purposes of data sharing, organizations could circumvent the limits imposed by this standard. Under the user expectations standards we adopt, a corporate affiliate is not, in general, the same party as an organization.

3.2.1.2.3 Branding

In many cases, branding aligns with ordinary user expectations. Unrelated websites rarely share branding. In company ownership scenarios, prominent language like "Brand A, provided by Company B" may be sufficient for the average user to understand that Brand A is owned by Company B and information shared with Brand A may also be shared with Company B.

But, in some cases, branding does not align with user expectations. Suppose Example Search owns a video sharing website, Example Video. Most users are aware that Example Video is a subsidiary of Example Search, and that the Example Video website differs from the Example Search website for historical reasons. The Example Video home page does not, however, include any branding reference to Example Search. Under a branding test, Example Search and Example Branding would be different parties. The user expectations test allows for factors, other than branding, that influence user understanding.

Branding may also fall short in informing user expectations. If most users have never heard of Company B, language like "Brand A, provided by Company B" may not be adequate for the average user to understand the relationship between Brand A and Company B. A user expectations test recognizes there may be instances where even conspicuous branding is inadequate to inform users.

3.2.2 Option 2: Discoverable Affiliates

3.2.2.1 Definition

A party is any commercial, nonprofit, or governmental organization, a subsidiary or unit of such an organization, or a person. For unique corporate entities to qualify as a common party with respect to this standard, those entities MUST be commonly owned and commonly controlled, and MUST make their parent affiliation (if any) easy discoverable to users.

3.2.2.2 Discussion

This section is non-normative.

This may be accomplished in many ways, including but not limited to, prominent and common branding on site pages, "one click away" within Privacy Policies, and, if available, a programmatic list of domains that share common ownership (affiliation).

3.3 First Parties and Third Parties

3.3.1 Option 1: Meaningful User Interaction

3.3.1.1 Definitions

A "first party" is any party, in a specific network interaction, that can infer with high probability that the user knowingly and intentionally communicated with it. Otherwise, a party is a third party.

A "third party" is any party, in a specific network interaction, that cannot infer with high probability that the user knowingly and intentionally communicated with it.

Note

Some participants are concerned this proposal is not implementable as-is.

3.3.1.2 Discussion

This section is non-normative.

3.3.1.2.1 Overview

We draw a distinction between those parties an ordinary user would or would not expect to share information with, "first parties" and "third parties" respectively. The delineation exists for three reasons.

First, when a user expects to share information with a party, she can often exercise control over the information flow. Take, for example, Example Social, a popular social network. The user may decide she does not like Example Social's privacy or security practices, so she does not visit examplesocial.com. But if Example Social provides a

social sharing widget embedded in another website, the user may be unaware she is giving information to Example Social and unable to exercise control over the information flow.

Second, we recognize that market pressures are an important factor in encouraging good privacy and security practices. If users do not expect that they will share information with an organization, it is unlikely to experience market pressure from users to protect the security and privacy of their information. In practice, moreover, third parties may not experience sufficient market pressure from first parties since increasingly third parties do not have a direct business relationship with the first party websites they appear on. We therefore require a greater degree of user control over information sharing with such organizations.

Last, third parties are often in a position to collect a sizable proportion of a user's browsing history – information that can be uniquely sensitive and easily associated with a user's identity. We wish to provide user control over such information flows.

We recognize that, unlike with a bright-line rule, there can be close calls in applying our standard for what constitutes a first party or a third party. But we believe that in practice, such close calls will be rare. The overwhelming majority of content on the web can be classified as first party or third party, with few cases of ambiguity in practice.

We require a confidence at a "high probability" before a party can consider itself a first party. Where there is reasonable ambiguity about whether a user has intentionally interacted with a party, it must consider itself a third party. Our rationale is that, in the rare close cases, a website is in the best position to understand its users' expectations. We therefore impose the burden of understanding user expectations on the website. We also wish, in close cases, to err on the side of conforming to user expectations and protecting user privacy. If the standard is insufficiently protective, ordinary users have limited recourse; if the standard imposes excessive limits, websites retain the safety valve of explicitly asking for user permission.

3.3.1.2.1.1 Common Examples and Use Cases

1. A user accesses an Example News article. The page includes an advertisement slot, which loads content from many companies other than Example News. Those companies are third parties.
2. A user accesses an Example News article. The page includes an analytics script that is hosted by Example Analytics, an analytics service. Example Analytics is a third party.
3. A user accesses an Example News article. It includes a social sharing widget from Example Social, a popular social network. Example Social is a third party.
4. A user visits Example Diary, which is hosted by the free blogging service Example Blog Hosting but located at examplediary.com. Example Blog Hosting is a third party.
5. A user launches Example Application, an app on a mobile device. The app includes a library from Example Advertising Network that displays ads. Example Advertising Network is a third party.

3.3.1.2.2 Multiple First Parties

Note

While this is not closed the idea there may be multiple first parties on one webpage based on meaningful interaction has little controversy, with some discussion around the margins about details of co-run sites.

There will almost always be only one party that the average user would expect to communicate with: the provider of the website the user has visited. But, in rare cases, users may expect that a website is provided by more than one party. For example, suppose Example Sports, a well known sports league, collaborates with Example Streaming, a well known streaming video website, to provide content at www.examplesportsonexamplestreaming.com. The website is prominently advertised and branded as being provided by both Example Sports and Example Streaming. An ordinary user who visits the website may recognize that it is operated by both Example Sports and Example Streaming.

3.3.1.2.3 User Interaction with Third-Party Content

A party may start out as a third party but become a first party later on, after a user interacts with it. If content from a third party is embedded on a first party page, the third party may become an additional first party if it can infer with high probability that the average user knowingly and intentionally communicated with it. If a user merely moused over, closed, or muted third-party content, the party would not be able to draw such an inference.

3.3.1.2.3.1 Examples and Use Cases

Note

Little controversy, with some discussion of adding examples of co-run sites unresolved by publication. Some think branding covers this case; others want to add an additional example: A user visits a contest website hosted by two parties (clearly branded): Company X and Company Y. The page includes links to both parties legal information (Privacy Policy, Terms of Service) with explanation that both parties are first parties on this website.

Example: Example Weather offers an unbranded weather widget that is embedded into websites, including Example News. The widget contains small links to Example Weather's website and privacy policy. A user visits Example News and scrolls through the weekly forecast in the Example Weather widget.

Discussion: Example Weather is a third party. The user has interacted with Example

Weather's widget, but an ordinary user would not expect that scrolling through the widget involves communicating with Example News.

Example: Example Social, a popular social network, hosts a social sharing button that other websites can embed. The button is colored and styled in the same fashion as Example Social's website, contains descriptive text that is specific to Example Social, includes Example Social's logo, and very frequently appears on Example Social's website. Example News embeds the Example Social button, and a user clicks it.

Discussion: Example Social is a first party once the user clicks its embedded social sharing button. The average user would understand that by clicking the button she is communicating with Example Social.

3.3.2 Option 2: Discoverable Ownership and Affiliates

Note

This language merely rehashes the discoverable affiliate definition of "party" and does not provide guidance as to when a party is operating on a first-party or third-party basis. This language would be better incorporated as non-normative discussion under the Option 2: Discoverable Affiliates definition of party above.

3.3.2.1 Definitions

A First Party is the entity that owns the Web site or has Control over the Web site the consumer visits. A First Party also includes the owner of a widget, search box or similar service with which a consumer interacts, even if such First Party does not own or have Control over the Web site where the widget or services are displayed to the consumer.

A First Party includes Affiliates of that First Party, but only to the extent that the Affiliate is (1) an entity that Controls, is Controlled by, or us under common Control with, the First Party; or (2) an entity where the relationship to the First Party is clear to consumers through co-branding or similar means.

A First Party must make reasonable efforts to disclose, in a manner easily discoverable by Users, its ownership or Control of a site or service, such as through branding on the site or service, disclosures in the privacy policy, or terms of use linked to that site or service.

Control of an entity means that one entity (1) is under significant common ownership or operational control of the other entity, or (2) has the power to exercise a controlling influence over the management or policies of the other entity. In addition, for an entity to be under the Control of another entity and be treated as a First Party under this standard, the entity must also adhere to DNT standard in this specification.

3.4 Network Interaction

3.4.1 Definition

A "network interaction" is an HTTP request and response, or any other sequence of logically related network traffic.

3.4.2 Discussion

This section is non-normative.

Determination of a party's status is limited to a single interaction because a party's status may be affected by time, context, or any other factor that influences user expectations.

3.5 Transactional data

Transactional data is information about the user's interactions with various websites, services, or widgets which could be used to create a record of a user's system information, online communications, transactions and other activities, including websites visited, pages and ads viewed, purchases made, etc.

3.6 Data collection, retention, use, and sharing

Note

The following text consists of proposed text that is meant to address ISSUE-16. This language is currently being actively debated.

Issue

ISSUE-16 : What does it mean to collect data? (caching, logging, storage, retention, accumulation, profile etc.)

1. A party "collects" data if the data comes within its control.
2. A party "retains" data if data remains within a party's control.
3. A party "uses" data if the party processes the data for any purpose other than storage.
4. A party "shares" data if the party enables another party to collect the data.

The definitions of collection, retention, use, and sharing are drafted expansively so as to comprehensively cover a party's user-information practices. These definitions do not require a party's intent; a party may inadvertently collect, retain, use, or share data. The definition of collection includes information that a party did not cause to be transmitted, such as protocol headers.

3.7 Tracking

Note

We are still working through how, or if, to define tracking. Some suggest the phrase "cross-site tracking" only. We will need to ensure both final recommendations use the same terms in the same way.

3.7.1 Option 1: Non-first Party Identifiers

Note

Concerns with this section include undefined term "user data" plus as written, this may apply more broadly than the authors intended

Tracking is the collection or use of user data via either a unique identifier or a correlated set of data points being used to approximate a unique identifier, in a context other than "first party" as defined in this document. This includes:

1. a party collecting data across multiple websites, even if it is a first party in one or more (but not all) of the multiple contexts
2. a third party collecting data on a given website
3. a first party sharing user data collected from a DNT-on user with third parties "after the fact".

Examples of tracking use cases include:

- personalized advertising
- cross-site analytics or market research that has not been de-identified
- automatic preference sharing by social applications

3.7.2 Option 2: Cross-site or Over Time

Tracking is defined as following or identifying a user, user agent, or device across multiple visits to a site (time) or across multiple sites (space).

Mechanisms for performing tracking include but are not limited to:

- assigning a unique identifier to the user, user agent, or device such that it will be conveyed back to the server on future visits;
- personalizing references or referral information such that they will convey the user, user agent, or device identity to other sites;
- correlating data provided in the request with identifying data collected from past requests or obtained from a third party; or,
- combining data provided in the request with de-identified data collected or obtained from past requests in order to re-identify that data or otherwise associate it with the user, user agent, or device.

A preference of "Do Not Track" means that the user does not want tracking to be engaged for this request, including any mechanism for performing tracking, any use of data retained from prior tracking, and any retention or sharing of data from this request for the purpose of future tracking, beyond what is necessary to enable:

1. the limited permitted uses defined in this specification;
2. the first-party (and third-parties acting as the first-party) to provide the service intentionally requested by the user; and
3. other services for which the user has provided prior, specific, and informed consent.

3.7.3 Option 3: Silence

One proposal is not to define "tracking," but rather to list what is, or is not, required and allowed in order to comply with the recommendation.

3.8 Consent

The term “affirmative, informed consent” is used throughout this document. While this terminology may ultimately be modified, some options for explaining the underlying idea are presented below:

Issue

ISSUE-69 : Should the spec say anything about minimal notice? (ie. don't bury in a privacy policy)

3.8.1 Option 1

"Affirmative, Informed Consent to be Tracked" means consent given by an affirmative action such as clicking a consent box in response to a clear and prominent request to ignore a "Do Not Track" setting that is distinct and separate from any other notifications or requested permissions.

3.8.2 Option 2

"Affirmative, Informed Consent to be Tracked" has been obtained when a mechanism to provide for or facilitate the acquisition and storage of permission to ignore the header has been made available to the user and the user has meaningfully interacted with the mechanism in a way that makes clear her intent to grant this permission.

3.8.3 Silence

Note

The hope is that this option will ensure consistency with EU regulations; it may not unless notice is included.

No definition, other than explicitly leaving the definition of consent to local rules.

3.9 Meaningful Interaction

Note

Wording needs polish to ensure it works with accessibility issues, but other than minor edits this is agreed upon.

"Meaningful Interaction" with a widget or window initially presented on a third-party basis

means clicking on such content (except to stop, close, silence, or otherwise impair the rendering of such content) or otherwise affirmatively engaging with the content in a manner that would reasonably be interpreted to express an affirmative intention to interact with that party. A user merely moving her cursor across the widget or window does not constitute "meaningful interaction."

3.10 User

A user is an individual human. When user-agent software accesses online resources, whether or not the user understands or has specific knowledge of a particular request, that request is made "by" the user.

3.11 User Agent

This specification uses the term user agent to refer to any of the various client programs capable of initiating HTTP requests, including but not limited to browsers, spiders (web-based robots), command-line tools, native applications, and mobile apps [*HTTP11*].

4. Compliance with an expressed tracking preference

4.1 Compliance by a first party

Note

This section consists of proposed text that is meant to address ISSUE-17: Data use by 1st Party, ISSUE-30: Will Do Not Track apply to offline aggregating or selling of data?, ISSUE-54: Can first party provide targeting based on registration information even while sending DNT, ISSUE-59: Should the first party be informed about whether the user has sent a DNT header to third parties on their site?, and ISSUE-91: Might want prohibitions on first parties re-selling data to get around the intent of DNT, and are pending discussion and **[PENDING REVIEW]**.

Note

Additional text has been prepared for each of these five options. This text will be added once an option has been decided upon.

The following are distinct options that have been proposed by members of the group:

1. **Only share if (1):** If an operator of a first party domain stores a request to which a [DNT-ON] header is attached, that operator MUST NOT share information about that stored communication to a third party, outside of the permitted uses as defined in this standard or specific, user-granted exceptions granted.
2. **Only share if (2):** For those users who send the DNT signal and have not granted a site-specific exception to the first party, first parties must NOT share user-specific data with third parties [This will leave room for sharing with Agents/Service Providers/Vendors to the 1st party &emdash; as well as sharing aggregate and anonymous data with "others" (general reporting, for example).]
3. **Do not remember:** When DNT is enabled, a 1st party should treat each session

with a user as an entirely new session unless it has been given permission to store his information and use it again.

4. **Silence:** This standard imposes no requirements on first-party websites. A first-party website MAY take steps to protect user privacy in responding to a Do Not Track request.
5. **Do not circumvent:**
    1. Sharing of Data with a Third-Party Website: A first-party website MUST NOT share data to a third-party website that the third-party website could not collect itself under this standard. A first-party website MAY otherwise transfer data to a third-party website.
    2. Additional Voluntary Measures: A first-party website MAY take additional steps to protect user privacy in responding to a Do Not Track request.
        1. Example Voluntary Measures (Non-Normative)
    3. Transfer of Data from a First-Party Website: If a third-party website receives data from a first-party website, the data is subject to the same collection, retention, sharing, and use limitations under this standard as if the third-party website had collected the data itself.

4.2 Intermediary compliance

Note

This issue is being addressed in the [*TRACKING-DNT*] specification.

4.3 Compliance by a third party

Note

This section consists of proposed text that is meant to address ISSUE-19 and ISSUE-39 and is pending discussion and **[PENDING REVIEW]**.

4.3.1 General compliance

4.3.1.1 Option 1: Simple Formulation

If the operator of a third-party domain receives a communication to which a [DNT-ON] header is attached:

1. that operator MUST NOT collect, share, or use information related to that communication outside of the permitted uses as defined within this standard and any explicitly-granted exceptions, provided in accordance with the requirements of this standard;
2. that operator MUST NOT use information about previous communications in which the operator was a third party, outside of the explicitly expressed permitted uses as defined within this standard;
3. that operator [MUST NOT or SHOULD NOT] retain information about previous communications in which the operator was a third party, outside of the explicitly expressed permitted uses as defined within this standard.

4.3.1.2 Option 2: More Detailed Formulation

Note

The following consists of proposed text that is meant to address ISSUE-71 and is pending discussion and **[PENDING REVIEW]**.

Issue

ISSUE-71 : Does DNT also affect past collection or use of past collection of info?

1. When a third party receives a DNT signal, it MUST NOT relate additional data from that HTTP request to existing profiles associated with that user-agent that are based on data that the third party has previously collected across sites over time; this is except as allowed by permitted uses stated elsewhere in this specification
2. Three alternatives:
    1. Additionally, the entity MUST NOT use identifiers that it can determine were collected from the same user agent before the DNT signal was received, except as allowed by permitted uses, for as long as it continues to receive a DNT signal from that user-agent.
    2. A third party MUST NOT associate collected data with either previous or future user profiles. Any third party data collected under operational purpose permitted uses MUST NEVER be profiled independently or associated with previous or future user profiles.
    3. When a third party receives a DNT signal, it MUST NOT retain data from that HTTP request that could be associated with an existing profile, except as allowed by permitted uses stated elsewhere in this specification.
3. The entity MAY take additional steps with respect to previously collected DNXT data such as deleting data before its usual expiration. However, as DNT signal affects only HTTP request that it accompanies and may be modified by the user, it is not recommended that special deletion take place without some notice to user(s).

4.3.1.2.1 Examples and use cases

This section is non-normative.

1. User visits Site A, to which Ad Network B delivers advertisements. Ad Network B has accumulated transactional information about User from User's visits to Site A and other non-affiliated sites in the past. However, User now sends DNT signal with HTTP request during this session on Site A. Ad Network B cannot add information from current HTTP request from Site A session to any profile it maintains on User. Since it must not collect and any data from this session and relate it to previously collected data, Network B must regard and treat him like completely unknown user to them, absent any permitted uses or override from user.

2. Same as above scenario. Based on transactional information collected about

User's visits to non-affiliated sites in the past, Ad Network B has placed User into Technology Shopper Segment. Since Ad Network B must not recognize User during sessions in which User is sending DNT signal via that browser, it cannot deliver Technology Shopper advertisement to User's browser, absent obtaining override from user. Ad Network B may instead choose to deliver a random ad, an ad based on the context of Site A, or an ad based on general location based on IP address transmitted with HTTP.

4.3.2 Geolocation compliance by a third party

Note

This section may move into, or have implications for, the section 3 definitions.

If the operator of a third-party domain receives a communication to which a [DNT-ON] header is attached:

1. Geo-location information that is more granular than postal code is too granular. Geolocation data MUST NOT be used at any level more granular than postal code. Note that while the number of people living in a postal code varies from country to country, postal codes are extant world-wide.
2. If specific consent has been granted for the use of more granular location data, than that consent prevails.

4.3.2.1 Discussion

This section is non-normative.

It is acceptable to use data sent as part of this particular network interaction when composing a response to a [DNT-ON] request, but it is not acceptable to store that data any longer than needed to reply. For instance, it would be appropriate to use an IP address to guess which country a user is in, to avoid showing them an advertisement for products or services unavailable where they live.

When using request-specific information to compose a reply, some levels of detail may feel invasive to users, and may violate their expectations about Do Not Track. These sorts of detailed assessments should be avoided.

4.3.2.1.1 Examples

Reasonable behavior

A user visits you from an IP address which a general geo-IP database suggests is in the NYC area, where it is 6pm on a Friday. You choose to show an advertisement for theaters and restaurants in the area.

Invasive behavior

A user visits you from an IP address which suggests that they are in a particular

ZIP+4, which has a distinctive demographic profile. Their user-agent indicates that they are a Mac user, further narrowing their expected profile. You serve them an ad for business within a few blocks of them which specializes in items which their expected profile indicates they may enjoy.

In this example, even though the decision about which ad to serve was based exclusively on request specific information, but was still tailored to a highly-specific user profile. In particular, the estimation of a user's location to within a single ZIP+4 may make a user feel that they are being followed closely, even if the decision was made on the fly, and the information was only held ephemerally.

Issue

ISSUE-19 : Data collection / Data use (3rd party)

Issue

ISSUE-88 : different rules for impression of and interaction with 3rd-party ads/content

4.4 Usage-based Permitted Uses

This section outlines potential permitted uses based on necessary business use. For all of these permitted uses, the complying entity must make reasonable data minimization efforts to ensure that only the data necessary for the permitted use be retained.

Note

The following text consists of proposed text that is meant to address ISSUE-23 , ISSUE-24 , ISSUE-25 , ISSUE-31 , ISSUE-34 , and ISSUE-49 and is pending discussion and **[PENDING REVIEW]**.

Issue

Should we explicitly identify goals and use cases in order to evaluate these permitted uses?

4.4.1 Permitted uses for operational use of data

Note

This section consists of proposed text that is meant to address ISSUE-22 and is pending discussion and **[PENDING REVIEW]**. We have active discussions in this area.

Issue

ISSUE-22 : Still have "operational use" of data (auditing of where ads are shown, impression tracking, etc.)

4.4.1.1 Discussion

This section is non-normative.

In order to preserve certain common and important data usages, while still protecting consumer privacy concerns, it will be necessary to provide operational purpose permitted uses for necessary business activities when the DNT signal is on. There are several key categories of data collection and use that must remain intact such that web site operators who are (in the vast majority) offering their services free of charge in exchange for advertising on their properties. Proposed permitted uses include:

1. Frequency Capping - A form of historical tracking to ensure the number of times a user sees the same ad is kept to a minimum. Perhaps related, sequential ad rotation. *Withdrawn.*
2. Financial Logging - Ad impressions and clicks (and sometimes conversions) events are tied to financial transactions (this is how online advertising is billed) and therefore must be collected and stored for billing and auditing purposes.
3. 3rd Party Auditing - Online advertising is a billed event and there are concerns with accuracy in impression counting and quality of placement so 3rd party auditors provide an independent reporting service to advertisers and agencies so they can compare reporting for accuracy.
4. Security - From traditional security attacks to more elaborate fraudulent activity, ad networks must have the ability to log data about suspected bad actors to discern and filter their activities from legitimate transactions. This information is sometimes shared across 3rd parties in cooperatives to help reduce the daisy-chain effect of attacks across the ad ecosystem.
5. Contextual Content or Ad Serving: A third-party may collect and use information contained with the user agent string (including IP address and referrer url) to deliver content customized to that information.
6. Research / Market Analytics
7. Product Improvement, or, more narrowly, Debugging

Discussion is ongoing as to how to define these permitted uses and whether or not all should be included in an permitted uses list.

4.4.2 Permitted use for Outsourcing

Note

This section consists of proposed text that is meant to address ISSUE-23 , ISSUE-72 , ISSUE-73 and ISSUE-49 and are pending discussion and **[PENDING REVIEW]**.

Note

This section may be moved to the section titled "First Parties and Third Parties"

Issue

Issue

ISSUE-49 : Third party as first party - is a third party that collects data on behalf of the first party treated the same way as the first party?

4.4.2.1 Requirements

A third-party site MAY operate as a first-party site if all the following conditions hold:

1. the third party's data collection, retention, and use practices comply with at least the requirements for first-parties;
2. the data collected by the third party is available only to the first party, and the third party has no independent right to use the data;
3. the third party makes commitments that are consistent with compliance with this standard and they do so in a form that is legally enforceable (directly or indirectly) by the first party, individual users, and regulators; data retention by the third party must not survive the end of this legal enforceability;
4. the third party undertakes reasonable technical precautions to prevent collecting [retention of?] data that could be correlated across first parties.

4.4.2.2 Discussion

This section is non-normative.

4.4.2.2.1 Overview

The rationale for rule (2) is that we allow the third party to stand in the first party's shoes – but go no further. The third party may not use the data it collects for "product improvement," "aggregate analytics," or any other purpose except to fulfill a request by a first party, where the results are shared only with the first party.

Rule (3) allows for the possibility of more than one level of outsourcing.

In rule (4), one component of reasonable technical precautions will often be using the same-origin policy to segregate information for each first-party customer.

Note that any data collected by the third party that is used, or may be used, in any way by any party other than the first party, is subject to the requirements for third parties.

4.4.2.2.2 Examples and use cases

ExampleAnalytics collects analytic data for ExampleProducts Inc.. It operates a site under the DNS analytics.exampleproducts.com. It collects and analyzes data on visits to ExampleProducts, and provides that data solely to ExampleProducts, and does not access or use it itself.

4.4.3 Permitted use for unidentifiable data

Note

This section consists of proposed text that is meant to address ISSUE-34 and is pending discussion and **[PENDING REVIEW]**.

Issue

ISSUE-34 : Possible permitted use for aggregate analytics

4.4.3.1 Requirements

A third party MAY collect, retain, and use any information from a user or user agent that, with high probability, could not be used to:

1. identify or nearly identify a user or user agent; or
2. correlate the activities of a user or user agent across multiple network interactions.

4.4.3.2 Discussion

This section is non-normative.

4.4.3.2.1 Overview

Note

Clarification is needed with regard to what is meant by the following text

This permitted use (like all permitted uses) may not be combined with other permitted uses unless specifically allowed. A third party acting within the outsourcing permitted use, for example, may not make independent use of the data it has collected even though the use involves unidentifiable data.

A rule to the contrary would provide a perverse incentive for third parties to press all permitted uses to the limit and then use the collected data within this permitted use.

A potential 'safe harbor' under this clause could be to retain only aggregate counts, not per-transaction records.

4.4.3.2.2 Examples and use cases

1. A third-party advertising network records the fact that it displayed an ad.

2. A third-party analytics service counts the number of times a popular page was loaded.

4.4.4 Other issues raised around permitted uses

Issue

ISSUE-24 : Possible permitted use for fraud detection and defense

Issue

ISSUE-25 : Possible permitted use for research purposes

Adherence to laws, legal and judicial process, regulations and so forth take precedence over this standard when applicable, but contractual obligations do not.

Issue

ISSUE-75 : How do companies claim permitted uses and is that technical or not?

Issue

ISSUE-31 : Minimization &emdash; to what extent will minimization be required for use of a particular permitted use? (conditional permitted uses)

Issue

ISSUE-92 : If data collection (even very specific with IP address, user agent, referrer) is time-limited, with very limited retention, is that still tracking?

Issue

ISSUE-89 : Does DNT mean at a high level: (a) no customization, users are seen for the first time, every time. (b) DNT is about data moving between sites.

Issue

ISSUE-97: Re-direction, shortened URLs, click analytics &emdash; what kind of tracking is this?

5. User-Granted Exceptions

Issue

ISSUE-66 : Can user be allowed to consent to both third party and first party to override general DNT?

Issue

ISSUE-93 : Should 1st parties be able to degrade a user experience or charge money for content based on DNT?

5.1 Introduction to user-granted exceptions

For the purposes of this document, a user-granted exception is a user-granted override of their default DNT status for one or more third parties within a given first party context.

It is possible for first parties to request, and users to set, user-granted exceptions to their default DNT status on a per-first party basis for the third parties that the first party works with. The goal of this is to allow first parties to communicate with their users about their options with respect to DNT within the context of that first party's web pages.

Note

Should Market Research be deemed a user-granted exception rather than a permitted use?

5.2 Opt-In to site-specific, user-granted exceptions

Note

The following consists of proposed text and is pending discussion and **[PENDING REVIEW]**.

When a DNT enabled user agent grants a site-specific, "user-granted" exception, the site places a site-specific opt-in mechanism on the user agent allowing the site to respond as a First Party. The DNT header must remain enabled so that if the user returns to the site, both the user's general preference for DNT and the site-specific, user-granted exception will be clear. When seeking a site-specific exception from the user, the site must describe to the user, via a direct link from the user-granted exception page, all purposes for which the tracking will be used.

5.3 Interaction with existing user privacy controls

As multiple systems may be setting, sending, and receiving DNT and/or Opt-Out signals at the same time, it'll be important to ensure industry and web browser vendors are on the same page with respect to honoring user choices in circumstances where "mixed signals" may be received.

As a general principle, more specific settings override less specific settings.

- No DNT Signal / No Opt-Out: Treat as DNT unset
- DNT Signal / No Opt-Out: Treat as DNT:1
- Opt-Out / No DNT Signal: Treat as DNT:1
- Opt-Out / DNT User-Granted Exception: Treat as DNT:0 for that site; DNT User-Granted Exception is honored

Issue

ISSUE-83 : How do you opt out if already opted in?

Issue

ISSUE-67 : Should opt-back-in be stored on the client side? [Not sure this doesn't belong in the technical spec]

5.4 Logged In

Issue

ISSUE-65 : How does logged in and logged out state work

5.4.1 Option 1: Logged In Honors DNT

If a user is logged into a first-party website and it receives a DNT:1 signal, the website MUST respect DNT:1 signal as a first party and SHOULD handle the user login as it normally would. If a user is logged into a third-party website, and the third party receives a DNT:1 signal, then it MUST respect the DNT:1 signal unless it falls under an exemption described in this document.

Example use cases:

- A user with DNT:1 logs into a search service called "Searchy". Searchy also operates advertisements on other websites. When the user is on a news website, Searchy receives DNT:1, and it must respect it, as Searchy is operating in a third-party context.
- A user with DNT:1 enabled visits a shopping website and logs in. The shopping website continues to provide recommendations, order history, etc. The shopping site includes third-party advertisements. Those third-parties continue to respect DNT:1. When the user purchases the items in their basket, a third-party financial transaction service is used. The user interacts with the third-party service, at which point it becomes first-party and may use previously collected data.
- A user with DNT:1 visits a website (Website A) that uses a third-party authentication service called "LogMeIn". The user logs into the site with his LogMeIn credentials. The user has interacted with LogMeIn, and now it can act as a first-party. Now the user vists Website B, which also uses the LogMeIn service, but is branded differently than Website A. LogMeIn MUST respect the DNT:1 signal until the user chooses to interact with LogMeIn in order to log into Website B.

5.4.2 Option 2: Silence on Logged In

No text on this topic at all, and let the existing rules work it out.

-

5.5 Enforcement/Compliance

5.5.1 Requirement

Note

Final wording awaits how the response is designed in the [*TRACKING-DNT*] recommendation, but we agree upon the general direction below.

In order to be in compliance with this specification, a party MUST make a public commitment that it complies with this standard.

5.5.2 Discussion

This section is non-normative.

A "public commitment" may consist of a statement in a privacy policy, a response header, a machine-readable tracking status resource at a well-known location, or any other reasonable means. This standard does not require a specific form of "public commitment."

Issue

ISSUE-21 : Enable external audit of DNT compliance

Note

We have reviewed one audit proposal that we declined to adopt. We may include a smaller-scoped proposal in the future, or may drop auditing all together.

A. Acknowledgements

This specification consists of input from many discussions within and around the W3C Tracking Protection Working Group, along with written contributions from Haakon Flage Bratsberg (Opera Software), Amy Colando (Microsoft Corporation), Roy T. Fielding (Adobe), Tom Lowenthal (Mozilla), Ted Leung (The Walt Disney Company), Jonathan Mayer (Stanford University), Ninja Marnau (Invited Expert), Matthias Schunter (IBM), John M. Simpson (Invited Expert), Kevin G. Smith (Adobe), Rob van Eijk (Invited Expert), Rigo Wenning (W3C), and Shane Wiley (Yahoo!).

The DNT header field is based on the original *Do Not Track* submission by Jonathan Mayer (Stanford), Arvind Narayanan (Stanford), and Sid Stamm (Mozilla). The DOM API for `NavigatorDoNotTrack` is based on the *Web Tracking Protection* submission by Andy Zeigler, Adrian Bateman, and Eliot Graff (Microsoft). Many thanks to Robin Berjon for ReSpec.js.

B. References

B.1 Normative references

[HTTP11]
R. Fielding; et al. *Hypertext Transfer Protocol - HTTP/1.1.* June 1999. Internet RFC 2616. URL: http://www.ietf.org/rfc/rfc2616.txt

[TRACKING-DNT]
Roy T. Fielding; David Singer. *Tracking Preference Expression (DNT).* 13 March 2012. W3C Working Draft. (Work in progress.) URL: http://www.w3.org/TR/2011/WD-tracking-dnt-20120313/

B.2 Informative references

[KnowPrivacy]
Joshua Gomez; Travis Pinnick; Ashkan Soltani. *KnowPrivacy.* 1 June 2009. URL: http://www.knowprivacy.org/report/KnowPrivacy_Final_Report.pdf

Tracking Preference Expression (DNT)

W3C Editor's Draft 04 June 2012

This version:
http://www.w3.org/2011/tracking-protection/drafts/tracking-dnt.html
Latest published version:
http://www.w3.org/TR/tracking-dnt/
Latest editor's draft:
http://www.w3.org/2011/tracking-protection/drafts/tracking-dnt.html
Previous version:
http://www.w3.org/TR/2012/WD-tracking-dnt-20120313/
Editors:
Roy T. Fielding, Adobe
David Singer, Apple

Abstract

This specification defines the technical mechanisms for expressing a tracking preference via the DNT request header field in HTTP, via an HTML DOM property readable by embedded scripts, and via properties accessible to various user agent plug-in or extension APIs. It also defines mechanisms for sites to signal whether and how they honor this preference, both in the form of a machine-readable tracking status resource at a well-known location and via a "Tk" response header field, and a mechanism for allowing the user to approve site-specific exceptions to DNT as desired.

Status of This Document

This section describes the status of this document at the time of its publication. Other documents may supersede this document. A list of current W3C publications and the latest revision of this technical report can be found in the W3C technical reports index at http://www.w3.org/TR/.

This document is a snapshot of live discussions within the Tracking Protection Working Group. It does not yet capture all of our work. For example, we have issues that are [PENDING REVIEW] with complete text proposals that did not make it into this draft. Text in white is typically [CLOSED]: we have reached a consensus decision. Text in blue boxes presents multiple options the group is considering. In some cases we are close to agreement, and in others we have more to discuss. An issue tracking system is available for recording raised, open, pending review, closed, and postponed issues regarding this document.

This document was published by the Tracking Protection Working Group as an Editor's Draft. If you wish to make comments regarding this document, please send them to public-tracking@w3.org (subscribe, archives). All feedback is welcome.

Publication as an Editor's Draft does not imply endorsement by the W3C Membership. This is a draft document and may be updated, replaced or obsoleted by other documents at any time. It is inappropriate to cite this document as other than work in progress.

This document was produced by a group operating under the 5 February 2004 W3C Patent Policy. W3C maintains a public list of any patent disclosures made in connection with the deliverables of the group; that page also includes instructions for disclosing a patent. An individual who has actual knowledge of a patent which the individual believes contains Essential Claim(s) must disclose the information in accordance with section 6 of the W3C Patent Policy.

Table of Contents

1. Introduction

The World Wide Web (WWW, or Web) consists of millions of sites interconnected through the use of hypertext. Hypertext provides a simple, page-oriented view of a wide variety of information that can be traversed by selecting links, manipulating controls, and supplying data via forms and search dialogs. A Web page is usually composed of many different information sources beyond the initial resource request, including embedded references to stylesheets, inline images, javascript, and other elements that might be automatically requested as part of the rendering or behavioral processing defined for that page.

Each of the hypertext actions and each of the embedded resource references might refer to any site on the Web, leading to a seamless interaction with the user even though the pages might be composed of information requested from many different and possibly independent Web sites. From the user's perspective, they are simply visiting and interacting with a single brand — the **first-party** Web property — and all of the technical details and protocol mechanisms that are used to compose a page representing that brand are hidden behind the scenes.

It has become common for Web site owners to collect data regarding the usage of their sites for a variety of purposes, including what led the user to visit their site (referrals), how effective the user experience is within the site (web analytics), and the nature of who is using their site (audience segmentation). In some cases, the data collected is used to dynamically adapt the content (personalization) or the advertising presented to the user (targeted advertising). Data collection can occur both at the first-party site and via third-party providers through the insertion of tracking elements on each page. A survey of these techniques and their privacy implications can be found in [*KnowPrivacy*].

People have the right to know how data about them will be collected and how it will be used. Empowered with that knowledge, individuals can decide whether to allow their online activities to be tracked and data about them to be collected. Many Internet companies use data gathered about people's online activities to personalize content and target advertising based on their perceived interests. While some people appreciate this personalization of content and ads in certain contexts, others are troubled by what they perceive as an invasion of their privacy. For them, the benefit of personalization is not worth their concerns about allowing entities with whom they have no direct relationship to amass detailed profiles about their activities.

Therefore, users need a mechanism to express their own preference regarding tracking that is both simple to configure and efficient when implemented. In turn, Web sites that are unwilling or unable to offer content without such targeted advertising or data collection need a mechanism to indicate those requirements to the user and allow them (or their user agent) to make an individual choice regarding exceptions.

This specification defines the HTTP request header field DNT for expressing a tracking preference on the Web, a well-known location (URI) for providing a machine-readable tracking status resource that describes a service's DNT compliance, the HTTP response header field Tk for resources to communicate their compliance or non-compliance with the user's expressed preference, and JavaScript APIs for determining DNT status and requesting a user-granted exception.

A companion document, [*TRACKING-COMPLIANCE*], more precisely defines the terminology of tracking preferences, the scope of its applicability, and the requirements on compliant first-party and third-party participants when an indication of tracking preference is received.

Issue

ISSUE-136: Resolve dependencies of the TPE on the compliance specification.
The WG has not come to consensus regarding the definition of tracking and the scope of DNT. As such, a site cannot actually say with any confidence whether or not it is tracking, let alone describe the finer details in a tracking status resource. This issue will be resolved by progress on the TCS document, though its resolution is a necessary prerequisite to understanding and correctly implementing the protocol defined by this document.

2. Notational Conventions

2.1 Requirements

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

2.2 Formal Syntax

This specification uses Augmented Backus-Naur Form [*ABNF*] to define network protocol syntax and WebIDL [*WEBIDL*] for defining scripting APIs.

2.3 Terminology

This specification uses the term **user agent** to refer to any of the various client programs capable of initiating HTTP requests, including, but not limited to, browsers, spiders (web-based robots), command-line tools, native applications, and mobile apps [*HTTP11*].

Note

The term **permitted use** is used to indicate a restricted set of conditions under which tracking is allowed in spite of the user's DNT preference. The term **user-granted exception** is used when the user has permitted tracking, usually in the form of a site-specific exception, for a given third-party. In general: permitted uses are additional permissions granted by the standard; user-granted exceptions are additional permissions granted by the user. These words are often confused when drafting new text.

3. Determining User Preference

The goal of this protocol is to allow a user to express their personal preference regarding tracking to each server and web application that they communicate with via HTTP, thereby allowing each service to either adjust their behavior to meet the user's expectations or reach a separate agreement with the user to satisfy all parties.

Key to that notion of expression is that it MUST reflect the user's choice, not the choice of some vendor, institution, or network-imposed mechanism outside the user's control. The basic principle is that a tracking preference expression is only transmitted when it reflects a deliberate choice by the user. In the absence of user choice, there is no tracking preference expressed.

A user agent MUST offer users a minimum of two alternative choices for a "Do Not Track" preference: `unset` or `on`. A user agent MAY offer a third alternative choice: `off`. If the user's choice is `on` or `off`, the tracking preference is **enabled**; otherwise, the tracking preference is **not enabled**.

A user agent MUST have a default tracking preference of `unset` (not enabled) unless a specific tracking preference is implied by the decision to use that agent. For example, use of a general-purpose browser would not imply a tracking preference when invoked normally as "SuperFred", but might imply a preference if invoked as "SuperDoNotTrack" or "UltraPrivacyFred". Likewise, a user agent extension or add-on MUST NOT alter the tracking preference unless the act of installing and enabling that extension or add-on is an explicit choice by the user for that tracking preference.

We do not specify how tracking preference choices are offered to the user or how the preference is enabled: each implementation is responsible for determining the user experience by which a tracking preference is enabled. For example, a user might select a check-box in their user agent's configuration, install an extension or add-on that is specifically designed to add a tracking preference expression, or make a choice for privacy that then implicitly includes a tracking preference (e.g., "Privacy settings: high"). Likewise, a user might install or configure a proxy to add the expression to their own outgoing requests.

Although some controlled network environments, such as public access terminals or managed corporate intranets, might impose restrictions on the use or configuration of installed user agents, such that a user might only have access to user agents with a predetermined preference enabled, the user is at least able to choose whether to make use of those user agents. In contrast, if a user brings their own Web-enabled device to a library or cafe with wireless Internet access, the expectation will be that their chosen user agent and personal preferences regarding Web site behavior will not be altered by the network environment, aside from blanket limitations on what resources can or cannot be accessed through that network. Implementations of HTTP that are not under control of the user MUST NOT express a tracking preference on their behalf.

4. Expressing a Tracking Preference

4.1 Expression Format

When a user has enabled a tracking preference, that preference needs to be expressed to all mechanisms that might perform or initiate tracking by third parties, including sites that the user agent communicates with via HTTP, scripts that can extend behavior on pages, and plug-ins or extensions that might be installed and activated for various media types.

When enabled, a tracking preference is expressed as either:

DNT	meaning
1	This user prefers not to be tracked on the target site.
0	This user prefers to allow tracking on the target site.

If a tracking preference is not enabled, then no preference is expressed by this protocol. This means that no expression is sent for each of the following cases:

- the user agent does not implement this protocol;
- the user has not yet made a choice for a specific preference; or,
- the user has chosen not to indicate a preference.

In the absence of regulatory, legal, or other requirements, servers MAY interpret the lack of an expressed tracking preference as they find most appropriate for the given user, particularly when considered in light of the user's privacy expectations and cultural circumstances. Likewise, servers might make use of other preference information outside the scope of this protocol, such as site-specific user preferences or third-party registration services, to inform or adjust their behavior when no explicit preference is expressed via this protocol.

Issue

ISSUE-59: Should the first party be informed about whether the user has sent a DNT header to third parties on their site?

Issue

ISSUE-111: Different DNT value to signify existence of site-specific exception (also linked to 4.1 and 6 below)

4.2 DNT Header Field for HTTP Requests

The **DNT** header field is hereby defined as the means for expressing a user's tracking preference via HTTP [*HTTP11*].

```
DNT-field-name  = "DNT"                          ; case-insensitive
DNT-field-value = ( "0" / "1" ) *DNT-extension   ; case-sensitive
DNT-extension   = %x21 / %x23-2B / %x2D-5B / %x5D-7E
                ; excludes CTL, SP, DQUOTE, comma, backslash
```

A user agent MUST send the **DNT** header field on all HTTP requests if (and only if) a tracking preference is enabled. A user agent MUST NOT send the DNT header field if a tracking preference is not enabled.

The DNT field-value sent by a user agent MUST begin with the numeric character "1" (%x31) if a tracking preference is enabled, the preference is for no tracking, and there is not a site-specific exception for the origin server targeted by this request.

The DNT field-value sent by a user agent MUST begin with the numeric character "0" (%x30) if a tracking preference is enabled and the preference is to allow tracking in general or by specific exception for the origin server targeted by this request.

Example

```
GET /something/here HTTP/1.1
Host: example.com
DNT: 1
```

An HTTP intermediary MUST NOT add, delete, or modify the DNT header field in requests forwarded through that intermediary unless that intermediary has been specifically installed or configured to do so by the user making the requests. For example, an Internet Service Provider MUST NOT inject "DNT: 1" on behalf of all of their users who have not selected a choice.

The remainder of the DNT field-value after the initial character is reserved for future extensions. User agents that do not implement such extensions MUST NOT send DNT-extension characters in the DNT field-value. Servers that do not implement such extensions SHOULD ignore anything beyond the first character.

DNT extensions are to be interpreted as modifiers to the main preference expressed by the first digit, such that the main preference will be obeyed if the recipient does not understand the extension. Hence, a DNT-field-value of "1xyz" can be thought of as "do not track, but if you understand the refinements defined by x, y, or z, then adjust my preferences according to those refinements." DNT extensions can only be transmitted when a tracking preference is enabled.

The extension syntax is restricted to visible ASCII characters that can be parsed as a single word in HTTP and safely embedded in a JSON string without further encoding (section 5.2.2 Representation). Since the DNT header field is intended to be sent on every request, when enabled, designers of future extensions ought to use as few extension characters as possible.

Note

This document does not have any implied or specified behavior for the user-agent treatment of cookies when DNT is enabled.

Issue

ISSUE-111: Different DNT value to signify existence of site-specific exceptions
Should the user agent send a different DNT value to a first party site if there exist site-specific exceptions for that first party? (e.g. DNT:2 implies "I have Do Not Track enabled but grant permissions to some third parties while browsing this domain") **[OPEN]**

4.3 JavaScript API to Detect Preference

4.3.1 Interface

The `NavigatorDoNotTrack` interface provides a means for the user's tracking preference to be expressed to web applications running within a page rendered by the user agent.

WebIDL

```
[NoInterfaceObject]
interface NavigatorDoNotTrack {
    readonly attribute DOMString doNotTrack;
};
```

4.3.2 Attributes

`doNotTrack` of type DOMString, readonly
: When a tracking preference is enabled, the doNotTrack attribute MUST have a string value that is the same as the DNT-field-value defined in section 4.2 DNT Header Field for HTTP Requests. If a tracking preference is not enabled, the value is `null`.

4.3.3 Implements

```
WebIDL

Navigator implements NavigatorDoNotTrack;
```

Objects implementing the `Navigator` interface [*NAVIGATOR*] (e.g., the `window.navigator` object) MUST also implement the `NavigatorDoNotTrack` interface. An instance of `NavigatorDoNotTrack` is obtained by using binding-specific casting methods on an instance of `Navigator`.

Issue

ISSUE-84: Make DNT status available to JavaScript
[OPEN] The API above has been deemed inadequate due to origin restrictions on embedded javascript by reference. We are awaiting new text to resolve this issue.

Issue

ISSUE-116: How can we build a JS DOM property which doesn't allow inline JS to receive mixed signals?

4.4 Plug-In APIs

User agents often include user-installable component parts, commonly known as **plug-ins** or **browser extensions**, that are capable of making their own network requests. From the user's perspective, these components are considered part of the user agent and thus ought to respect the user's configuration of a tracking preference. However, plug-ins do not normally have read access to the browser configuration.

Note

It is unclear whether we need to standardize the plug-in APIs or if we should rely on it being defined per user agent based on general advice here. No plug-in APIs have been proposed yet.

4.5 Tracking Preference Expressed in Other Protocols

A user's tracking preference is intended to apply in general, regardless of the protocols being used for Internet communication. The protocol expressed here is specific to HTTP communication; however, the semantics are not restricted to use in HTTP; the same semantics may be carried by other protocols, either in future revisions of this specification, or in other specifications.

When it is known that the user's preference is for no tracking, compliant services are still required to honor that preference, even if other protocols are used. For example, re-directing to another protocol in order to avoid receipt of the header is not compliant.

Note

The last paragraph may be more appropriate in the compliance document, as it discusses compliance.

5. Communicating a Tracking Status

5.1 Overview

The Tracking Protection protocol is designed to be applicable regardless of the response from servers that receive the tracking preference expression, allowing conformance to be achieved without impacting the operational performance of site resources. However, there is also a desire to support verification or pre-flight testing of a site's conformance with this protocol for evaluating conformance before sending data, enabling specialized user interfaces, discovering the scope of protocol deployment, and testing adherence to potential regulations.

This section explains how a user agent MAY discover an origin server's tracking status for a given resource. It defines a REQUIRED well-known tracking status resource for describing a machine-readable tracking status and a Tk response header field that MAY be sent in any HTTP response and MUST be sent in responses to requests that modify the tracking status for that user agent.

5.2 Tracking Status Resource

5.2.1 Definition

An origin server MUST provide a **tracking status resource** at the well-known identifier [*RFC5785*]

```
/.well-known/dnt
```

(relative to the URI of that origin server) for obtaining information about the potential tracking behavior of resources provided by that origin server. A tracking status resource MAY be used for verification of DNT support, as described in section 5.2.4 Using the Tracking Status.

A valid retrieval request (e.g., a GET in HTTP) on the well-known URI MUST result in either a successful response containing a machine-readable representation of the site-wide tracking status, as defined below, or a sequence of redirects that leads to such a representation. A user agent MAY consider failure to provide access to such a representation equivalent to the origin server not implementing this protocol. The representation MAY be cached, as described in section 5.2.5 Caching.

If an origin server has multiple, resource-specific tracking policies, such that the tracking status might differ depending on some aspect of the request (e.g., method, target URI, header fields, data, etc.), the origin server MAY provide an additional subtree of well-known resources corresponding to each of those distinct tracking statuses. The Tk response header field (section 5.3 Tk Header Field for HTTP Responses) can include a status-id to indicate which specific tracking status resource applies to the current request. This subtree of resources is called the **tracking status resource space**.

The **tracking status resource space** is defined by the following URI Template [*URI-TEMPLATE*]:

```
/.well-known/dnt{/status-id}
```

where the value of `status-id` is a string of URI-safe characters provided by a Tk field-value in response to a prior request. For example, a prior response containing

```
Tk: 1;fRx42
```

refers to the specific tracking status resource

```
/.well-known/dnt/fRx42
```

Resources within the tracking status resource space are represented using the same format as a site-wide tracking status resource.

When sending a request for the tracking status, a user agent SHOULD include any cookie data [*COOKIES*] (set prior to the request) that would be sent in a normal request to that origin server, since that data might be needed by the server to determine the current tracking status. For example, the cookie data might indicate a prior out-of-band decision by the user to opt-out or consent to tracking by that origin server.

All requests on the tracking status resource space, including the site-wide tracking status resource, MUST NOT be tracked, irrespective of the presence, value, or absence of a DNT header field, cookies, or any other information in the request. In addition, all responses to those requests, including the responses to redirected tracking status requests, MUST NOT have Set-Cookie or Set-Cookie2 header fields and MUST NOT have content that initiates tracking beyond what was already present in the request. A user agent SHOULD ignore, or treat as an error, any Set-Cookie or Set-Cookie2 header field received in such a response.

5.2.2 Representation

The representation of a tracking status resource SHALL be provided in the "application/json" format [*RFC4627*] and MUST conform to the ABNF in section 5.2.6 Status-object ABNF. The following is an example tracking status representation that illustrates all of the fields defined by this specification, most of which are optional.

Example

```
{
  "tracking": true,
  "received": "1",
  "response": "t1",
  "same-party": [
    "example.com",
    "example_vids.net",
    "example_stats.com"
  ],
  "partners": [
    "api.example-third-party.com"
  ],
  "policy": "/tracking.html",
  "control": "http://example-third-party.com/your/data"
}
```

A tracking status representation consists of a single status-object containing members that describe the tracking status applicable to this user agent's request.

A status-object MUST have a member named `tracking` with a boolean value. A value of `false` indicates that the corresponding resources do not perform tracking as it is defined by [*TRACKING-COMPLIANCE*]. A value of `true` indicates that the corresponding resource performs tracking and claims to conform to all tracking compliance requirements applicable to this site.

For example, the following demonstrates a minimal tracking status representation that is applicable to any resource that does not perform tracking.

Example

```
{"tracking": false}
```

If `tracking` is `true`, the status-object MUST include two additional members, named `received` and `response`, and MAY include other members as described below.

The `received` member MUST have either a string value equal to the DNT-field-value received in that request or the value `null` if no DNT-field-value was received. Any invalid characters received in the DNT-field-value MUST be elided from the string value to ensure that invalid data is not injected into the JSON result.

The `response` member MUST have a string value that indicates the status of tracking applicable specifically to this user in light of the received DNT-field-value. The string value begins with `t` (tracking), `n` (not tracking), or `s` (see the more specific tracking status resource), and MAY be followed by alphanumeric characters that indicate qualifiers for that status. The defined qualifier characters and their meanings are described in section Not foundstatus-reponse-value.

An OPTIONAL member named `same-party` MAY be provided with an array value containing a list of domain names that the origin server claims are the same party, to the extent they are referenced on this site, since all data collected via those references share the same data controller.

An OPTIONAL member named `partners` MAY be provided with an array value containing a list of domain names for third-party services that might be invoked while using this site but do not share the same data controller as this site.

An OPTIONAL member named `policy` MAY be provided with a string value containing a URI-reference to a human-readable document that describes the tracking policy for this site. The content of such a policy document is beyond the scope of this protocol and only supplemental to what is described by this machine-readable tracking status representation.

An OPTIONAL member named `control` MAY be provided with a string value containing a URI-reference to a resource for giving the user control over personal data collected by this site. Such control might include the ability to review past data collected, delete some or all of the data, provide

additional data (if desired), or "opt-in", "opt-out", or otherwise modify an out-of-band consent status regarding data collection by this site. The design of such a resource, the extent to which it can provide access to that data, and how one might implement an out-of-band consent mechanism is beyond the scope of this protocol.

Additional `extension` members MAY be provided in the status-object to support future enhancements to this protocol. A user agent SHOULD ignore extension members that it does not recognize.

Note that the tracking status resource space applies equally to both first-party and third-party services. An example of a third-party tracking status is

Example

```
{
  "tracking": true,
  "received": "1",
  "response": "n",
  "policy": "/privacy.html",
  "control": "/your/data",
}
```

Issue

ISSUE-47: Should the response from the server indicate a policy that describes the DNT practices of the server?
[PENDING REVIEW] The tracking status resource is a machine-readable policy and provides a mechanism for supplying a link to a human-readable policy.

Issue

ISSUE-61: A site could publish a list of the other domains that are associated with them
[PENDING REVIEW] The same-party and partners members provide a means to list first-party and third-party domains, respectively.

Issue

ISSUE-124: Alternative DNT implementations that replace HTTP headers with something else
[PENDING REVIEW] The tracking status resource minimizes bandwidth usage because only a small proportion of user agents are expected to perform active verification, status would only be requested once per site per day, and the response can be extensively cached.

5.2.3 Response Value

When present, the tracking status response member's value consists of a string of characters that starts with the tracking status, signified by `t` (tracking), `n` (not tracking), or `s` (see the more specific tracking status resource), and MAY be followed by a set of qualifier characters indicating reasons or limitations applicable to that status. Multiple qualifiers can be provided.

qualifier	meaning
1	First-party: The origin server acts as a first-party for requests on this resource, either in all contexts when no "3" qualifier is present or only for the domains listed in same-party.
3	Third-party: The origin server acts as a third-party for requests on this resource, either in all contexts when no "1" qualifier is present or only for the domains not listed in same-party.
a	Audit: Tracking is limited to that necessary for an external audit of the service context and the data collected is minimized accordingly.
c	Ad frequency capping: Tracking is limited to frequency capping and the data collected is minimized accordingly.
p	Prior consent: The origin server believes it has received prior explicit and informed consent for tracking this user, user agent, or device.
f	Fraud prevention: Tracking is limited to that necessary for preventing or investigating fraudulent behavior and security violations; the data collected is minimized accordingly.
l	Local constraints: Tracking is limited to what is required by local law, rule, or regulation and the data collected is minimized accordingly.
r	Referrals: Tracking is limited to collecting referral information and the data collected is minimized accordingly.

Qualifiers that indicate limitations on tracking correspond to the specific permitted uses in [*TRACKING-COMPLIANCE*]. An origin server indicating one or more of those permitted uses also indicates that it conforms to the requirements associated with those permitted uses. Multiple limitation qualifiers mean that multiple permitted uses of tracking might be present and that each such use conforms to the associated requirements. All limitation qualifiers imply some form of tracking might be used and thus MUST NOT be provided with a tracking status that begins with `n` (not tracking).

A `1` qualifier indicates that the resource has been designed for use within a first-party context and will conform to the requirements on tracking by a first-party. A `3` qualifier indicates that the resource has been designed for use within a third-party context and will conform to the requirements on tracking by a third-party. If both qualifiers are present, the resource is designed to dynamically adjust its tracking behavior according to the context in which it is used, and thus conforms to first-party requirements when used in a first-party context and third-party requirements when used in a third-party context.

A `p` qualifier indicates that the origin server believes it has obtained prior explicit and informed consent for tracking the requesting user agent, perhaps via some mechanism not defined by this specification, and that prior consent overrides the tracking preference expressed by this protocol. When prior consent is indicated, the tracking status object SHOULD include a `control` member that references a resource for modifying this consent.

Future extensions to this protocol might define additional characters as qualifiers from the `ext-qualifier` set (consisting of the remaining unused lowercase letters, dot, dash, and underscore). Recipients SHOULD ignore extension qualifiers that they do not understand.

Issue

ISSUE-136: Resolve dependencies of the TPE on the compliance specification.
The list of qualifiers is intended to correspond to constraints and permitted uses identified by [*TRACKING-COMPLIANCE*], and at some point might perhaps even move to that document in the sections defining the permitted uses. The above list will be updated accordingly.

Issue

ISSUE-137: Does hybrid tracking status need to distinguish between first party (1) and outsourcing service provider acting as a first party (s)
[PENDING REVIEW] No, a third party that satisfies the requirements for acting as a first party will communicate to users as the first party.

5.2.4 Using the Tracking Status

A key advantage of providing the tracking status at a resource separate from the site's normal services is that the status can be accessed and reviewed prior to making use of those services and prior to making requests on third-party resources referenced by those services. In addition, the presence (or absence) of a site-wide tracking status representation is a means for testing deployment of this standard and verifying that a site's claims regarding tracking are consistent with the site's observed behavior over time.

A user agent MAY check the tracking status for a given resource URI by making a retrieval request for the well-known address `/.well-known/dnt` relative to that URI.

If the response is an error, then the service does not implement this standard. If the response is a redirect, then follow the redirect to obtain the tracking status (up to some reasonable maximum of redirects to avoid misconfigured infinite request loops). If the response is successful, obtain the tracking status representation from the message payload, if possible, or consider it an error.

Once the tracking status representation is obtained, parse the representation as JSON to extract the Javascript status-object. If parsing results in a syntax error, the user agent SHOULD consider the site to be non-conformant with this protocol.

The status-object is supposed to have a member named `tracking` with a boolean value. If the value is "false", then no tracking is performed for the URI being checked. If the value is "true", then examine the member named `received` to verify that the DNT header field sent by the user agent has been correctly received by the server. If the `received` value is incorrect, there may be an intermediary interfering with transmission of the DNT request header field.

If the `received` value is correct, then examine the member named `response` to see what the origin server has claimed regarding the tracking status for this user agent in light of the received DNT-field-value.

If the first character of the `response` value is "n", then the origin server claims that it will not track the user agent for requests on the URI being checked for at least the next 24 hours or until the Cache-Control information indicates that this response expires, as described below.

If the first character of the `response` value is "t", then the origin server claims that it might track the user agent for requests on the URI being checked for at least the next 24 hours or until the Cache-Control information indicates that this response expires.

If the first character of the `response` value is "s", then the origin server has multiple tracking status representations and the specific one applicable to each request is indicated by a status-id within the Tk field-value of the corresponding response.

The remaining characters of the `response` value might indicate qualifiers for the above choices or limitations that the origin server will place on its tracking.

The others members of the status-object MAY be used to help the user agent differentiate between a site's first-party and third-party services, to provide links to additional human-readable information related to the tracking policy, and to provide links for control over past data collected or over some consent mechanism outside the scope of this protocol.

5.2.5 Caching

If the tracking status is applicable to all users, regardless of the received DNT-field-value or other data received via the request, then the response SHOULD be marked as cacheable and assigned a time-to-live (expiration or max-use) that is sufficient to enable shared caching but not greater than the earliest point at which the service's tracking behavior might increase. For example, if the tracking status response is set to expire in seven days, then the earliest point in time that the service's tracking behavior can be increased is seven days after the policy has been updated to reflect the new behavior, since old copies might persist in caches until the expiration is triggered. A service's tracking behavior can be reduced at any time, with or without a corresponding change to the tracking status resource.

If the tracking status is only applicable to all users that have the same "DNT-field-value", then either the response MUST include a Cache-Control header field with one of the directives "no-cache", "no-store", "must-revalidate", or "max-age=0", or the response MUST include a Vary header field that includes "DNT" in its field-value.

If the tracking status is only applicable to the specific user that requested it, then the response MUST include a Cache-Control header field with one of the directives "no-cache", "no-store", "must-revalidate", or "max-age=0".

Regardless of the cache-control settings, it is expected that user agents will check the tracking status of a service only once per session (at most). A public Internet site that intends to change its tracking status to increase tracking behavior MUST update the tracking status resource in accordance with that planned behavior at least twenty-four hours prior to activating that new behavior on the service.

A user agent that adjusts behavior based on active verification of tracking status, relying on cached tracking status responses to do so, SHOULD check responses to its state-changing requests (e.g., POST, PUT, DELETE, etc.) for a Tk header field with the update-needed field-value, as described in section 5.3.3 Indicating an Interactive Status Change.

5.2.6 Status-object ABNF

The representation of a site's machine-readable tracking status MUST conform to the following ABNF for status-object, except that the members within each member-list MAY be provided in any order.

```
status-object = begin-object member-list end-object
member-list   = tracking        ns tracking-v
                [ vs received   ns received-v   ]
                [ vs response   ns response-v   ]
                [ vs same-party ns same-party-v ]
                [ vs partners   ns partners-v   ]
                [ vs policy     ns policy-v     ]
                [ vs control    ns control-v    ]
                *( vs extension )

tracking      = %x22 "tracking" %x22
tracking-v    = true / false

received      = %x22 "received" %x22
received-v    = null / string

response      = %x22 "response" %x22
response-v    = %x22 r-codes %x22

r-codes       = (%x74 / %x6E / %x73) *qualifier

qualifier     = "1"   ; "1" — first-party
              / "3"   ; "3" — third-party
              / %x61  ; "a" — audit
              / %x63  ; "c" — ad frequency capping
              / %x66  ; "f" — fraud prevention
              / %x6C  ; "l" — local law, rule, or regulation
              / %x70  ; "p" — prior consent
              / %x72  ; "r" — referrals
              / ext-qualifier

ext-qualifier = %x2D-2E / "0" / "2" / %x34-39 / %x5F
              / %x62 / %x64-65 / %x67-6B / %x6D / %x6F
              / %x71 / %x75-7A

same-party    = %x22 "same-party" %x22
same-party-v  = array-of-strings

partners      = %x22 "partners" %x22
partners-v    = array-of-strings

policy        = %x22 "policy" %x22
policy-v      = string       ; URI-reference

control       = %x22 "control" %x22
control-v     = string       ; URI-reference

extension     = object

array-of-strings = begin-array
                   [ string *( vs string ) ]
                   end-array

ns            = <name-separator  (:), as defined in [RFC4627]>
vs            = <value-separator (,), as defined in [RFC4627]>

begin-array   = <begin-array     ([), as defined in [RFC4627]>
end-array     = <end-array       (]), as defined in [RFC4627]>
begin-object  = <begin-object    ({), as defined in [RFC4627]>
end-object    = <end-object      (}), as defined in [RFC4627]>
object        = <object, as defined in [RFC4627]>
string        = <string, as defined in [RFC4627]>
true          = <true,   as defined in [RFC4627]>
false         = <false,  as defined in [RFC4627]>
null          = <null,   as defined in [RFC4627]>
```

5.3 Tk Header Field for HTTP Responses

5.3.1 Definition

As a supplement to the tracking status resource, the **Tk** response header field is defined as an OPTIONAL means for indicating DNT conformance and as a REQUIRED means for indicating that a state-changing request has resulted in an interactive change to the tracking status for this user agent.

```
Tk-field-name   =  "Tk"     ; case-insensitive
Tk-field-value  =  tracking-design [ ";" status-id ]
tracking-design =  tracking-never
                        /  tracking-first
                        /  tracking-third
                        /  update-needed
tracking-never  =  "0"
tracking-first  =  "1"
tracking-third  =  "3"
update-needed   =  %x75     ; lowercase "u"
```

Issue

ISSUE-107: Exact format of the response header?
[PENDING REVIEW] See the proposal in this section.

5.3.2 Indicating Tracking Design

The Tk field-value begins with a single character tracking-design that indicates how the target resource conforms to [*TRACKING-COMPLIANCE*]. We refer to this as the tracking design because it reflects only how the resource is designed to work, rather than the current status of tracking for this requesting user agent or received DNT field-value. Separating the design and status allows conformance to this protocol to be indicated without having a negative impact on caching of responses.

An origin server MAY send a Tk header field in a response with a tracking-design of "0" to indicate that the resource never performs tracking as it is defined by [*TRACKING-COMPLIANCE*]. This has the same meaning as `{"tracking": "false"}` in the tracking status resource.

Example

```
Tk: 0
```

An origin server MAY send a Tk header field in a response with a tracking-design of "1" to indicate that the resource does perform tracking (though not necessarily for every request), conforms to [*TRACKING-COMPLIANCE*], and considers itself to be the first-party for this request.

Example

```
Tk: 1
```

An origin server MAY send a Tk header field in a response with a tracking-design of "3" to indicate that the resource does perform tracking (though not necessarily for every request), conforms to [*TRACKING-COMPLIANCE*], and considers itself to be a third-party for this request.

Example

```
Tk: 3
```

Issue

ISSUE-120: Should the response header be mandatory (MUST) or recommended (SHOULD)
[PENDING REVIEW] The site-wide resource is mandatory; the header field is optional, except for the single MUST case below.

5.3.3 Indicating an Interactive Status Change

We anticipate that interactive mechanisms might be used, beyond the scope of this specification, that have the effect of asking for and obtaining prior consent for tracking, or for modifying prior indications of consent. For example, the tracking status resource's status-object defines a `control` member that can refer to such a mechanism. Although such out-of-band mechanisms are not defined by this specification, their presence might influence the tracking status object's response value.

When an origin server provides a mechanism via HTTP for establishing or modifying out-of-band tracking preferences, the origin server MUST indicate within the mechanism's response when a state-changing request has resulted in a change to the tracking status for that server. This indication of an interactive status change is accomplished by sending a Tk header field in the response with a tracking-design of lowercase "u" (update-needed).

Example

```
Tk: u
```

5.3.4 Indicating a Specific Tracking Status Resource

If an origin server has multiple, resource-specific tracking policies, such that the tracking status might differ depending on some aspect of the request (e.g., method, target URI, header fields, data, etc.), the origin server MAY provide an additional subtree of well-known resources corresponding to each of those distinct tracking statuses. The OPTIONAL status-id portion of the Tk field-value indicates which specific tracking status resource applies to the current request.

For example, a response containing

```
Tk: 1;fRx42
```

indicates that the target resource conforms to this protocol as a first-party and the current tracking status can be obtained by performing a retrieval request on

```
/.well-known/dnt/fRx42
```

5.4 Status Code for Tracking Required

An HTTP error response status code might be useful for indicating that the site refuses service unless the user either logs into a subscription account or agrees to an exception to DNT for this site and its contracted third-party sites.

Issue

ISSUE-128: HTTP error status code to signal that tracking is required?

6. User-Granted Exceptions

Issue

ISSUE-111: Different DNT values to signify existence of site-specific exceptions

6.1 Overview

This section is non-normative.

User-granted exceptions to Do Not Track, including site-specific exceptions, are managed by the user agent. A resource should rely on the DNT header it receives to determine the user's preference for tracking with respect to that particular request. An API is provided so that sites may request and check the status of exceptions for tracking.

We anticipate that many user-agents might provide a prompt to users when this API is used, or to store exceptions. Questions of user interface specifics — for granting, configuring, storing, syncing and revoking exceptions — are explicitly left open to implementers.

6.2 Motivating principles and use cases

This section is non-normative.

The following principles guide the design of user-agent-managed exceptions.

- Content providers may wish to prompt visitors to their properties to “opt back in” to tracking for behavioral advertising or similar purposes when they arrive with the Do Not Track setting enabled.
- Privacy-conscious users may wish to view or edit all the exceptions they’ve granted in a single, consistent user interface, rather than managing preferences in a different way on every content provider or tracker’s privacy page.
- Granting an exception in one context (while browsing a news site) should not apply that exception to other contexts (browsing a medical site) that may not be expected.
- Tracking providers should not ever have to second-guess a user’s expressed Do Not Track preference.
- The solution should not require cross-domain communication between a first party publisher and its third parties.

When asking for a site-specific exception, the top-level domain making the request may be making some implicit or explicit claims as to the actions and behavior of its third parties; for this reason, it might want to establish exceptions for only those for which it is sure that those claims are true. (Consider a site that has some trusted advertisers and analytics providers, and some mashed-up content from less-trusted sites). For this reason, there is support both for explicitly named sites, as well as support for granting an exception to all third-parties on a given site (site-wide exception, using the wild-card "*").

There are some cases in which a user may desire a site to be allowed to track them on any top-level domain. An API is provided so that the site and the user may establish such a web-wide exception.

6.3 Exception model

6.3.1 Introduction

This API considers exceptions which are double-keyed to two domains: the **site**, and the **target**. A user might — for instance — want AnalytiCo to track them on Example News, but not on Example Medical. To simplify language used in this API specification, we define two terms:

- **Top-Level Domain (TLD)** is the domain name of the top-level document origin of this DOM: essentially the fully qualified domain name in the address bar. For all these APIs, this MUST match the script origin.
- A **target** site is a domain name which is the target of an HTTP request, and which may be an origin for embedded resources on **the indicated top-level domain**.

For instance, if the document at `http://web.exnews.com/news/story/2098373.html` references the resources `http://exnews.analytico.net/1x1.gif` and `http://widgets.exsocial.org/good-job-button.js`, **the top-level domain** is `web.exnews.com`; `exnews.analytico.net` and `widgets.exsocial.org` are both **targets**.

Issue

ISSUE-112: How are sub-domains handled for site-specific exceptions?
Should a request for a tracking exception apply to all subdomains of the first party making the request? Or should a first party explicitly list the subdomains that it's asking for? Similarly, should third party subdomains be allowed (e.g. *.tracker.com)?
Proposal: Exceptions are requested for fully-qualified domain names.

The domains that enter into the behavior of the APIs include:

- As described above, the **Top-Level Domain (TLD)**;
- The domain of the origin of the script;
- Domain names passed to the API;
- Domains declared in the well-known resource as 'partners'.

Note

Note that these strict, machine-discoverable, concepts may not match the definitions of first and third party; in particular, sites themselves need to determine (and signal) when they get 'promoted' to first party by virtue of user interaction; the UA will not change the DNT header it sends them.

During the execution of these APIs, the top-level browsing domain and the domain origin of the script MUST match, otherwise no action is taken, and an error value returned.

The calls causes the following steps to occur:

- First, the UA somehow confirms with the user that they agree to the grant of exception;
- If they agree, then the UA adds to its local database one or more site-pair duplets (top-level-domain, other-domain); one or other of these may be a wild-card ("*");
- While the user is browsing a given site [top-level domain], and a DNT header is to be sent to a target domain, if the duplet [top-level domain, target domain] matches any duplet in the database, then a DNT:0 header is sent, otherwise DNT:1 is sent.

6.3.2 Exception use by browsers

If a user wishes to be tracked by a **target** on the **top-level domain**, this should result in two user-agent behaviors:

1. If requests to the **target** for resources that are part of the DOM for pages on **top-level domain** include a DNT header, that header MUST be DNT:0.
2. Responses to the JavaScript API indicated should be consistent with this user preference (see below).

Issue

What is the effect of re-directs, when the source of the re-direct would get a different DNT header than the target, using these matching rules?

Note

It is left up to individual user-agent implementations how to determine and how and whether to store users' tracking preferences.

When such an explicit list of domains is provided through the API, their names might mean little to the user. The user might, for example, be told that such-and-such top-level domain is asking for an exception for a specific set of sites, rather than listing them by name.

Conversely, if a wild-card is used, the user may be told that the top-level domain is asking for an exception for all third-parties that are, or will be, embedded in it. The API might fetch the list of sites currently declared in the well-known URI as 'partners' as an example of the third-parties involved, but it should be noted that the partners list, and the set of embedded domains, might change after the API process is complete, and that the wild-card in the database applies dynamically to all sites that might be embedded, not just to the current 'partners' list.

Issue

ISSUE-111: Different DNT values to signify existence of user-granted exception
Should the user agent send a different DNT value to a first party site if there exist user-granted exceptions for that first party? (e.g. DNT:2 implies "I have Do Not Track enabled but grant permissions to some third parties while browsing this domain")
Proposal: No, this API provides client-side means for sites to request that information. Sites may also employ cookies to recall a user's past response. Finally, a site may add [self, self] to the database as part of its request, and it will then get DNT:0.

6.4 JavaScript API for site-specific exceptions

6.4.1 API to request site-specific exceptions

WebIDL

```
[NoInterfaceObject]
interface NavigatorDoNotTrack {
    void requestSiteSpecificTrackingException (sequence<DOMString> arrayOfDomainStrings, TrackingResponseCallback callback, optional
};
```

6.4.1.1 Methods

requestSiteSpecificTrackingException

Called by a page to request or confirm a user-granted tracking exception.

Parameter	Type	Nullable	Optional	Description
arrayOfDomainStrings	sequence<DOMString>	✘	✘	
callback	TrackingResponseCallback	✘	✘	
siteName		✘	✔	
explanationString		✘	✔	
detailURI		✘	✔	

Return type: void

WebIDL

```
[Callback, NoInterfaceObject]
interface TrackingResponseCallback {
    void handleEvent (boolean granted);
};
```

6.4.1.2 Methods

handleEvent

The callback is called by the user agent to indicate the user's response.

Parameter	Type	Nullable	Optional	Description
granted	boolean	✘	✘	

Return type: void

The requestSiteSpecificTrackingException method takes two mandatory arguments:

- `arrayOfDomainStrings`, a JavaScript array of strings, and
- `callback`, a method that will be called when the request is complete.

It also takes three optional arguments:

- `siteName`, a string for the name of the top-level domain (script origin),
- `explanationString`, a short explanation of the request, and
- `detailURI`, a location at which further information about this request can be found.

Each string in `arrayOfDomainStrings` specifies a **target**. The special string "*" signifies all **targets**. When called, `requestSiteSpecificTrackingException` MUST return immediately, then asynchronously determine whether the user grants the requested exceptions.

The `granted` parameter passed to the callback is the user's response; `true` indicates the user grants an exception on **top-level domain** for all of the **target**s specified in `arrayOfDomainStrings`. The response `false` indicates that the user does not want an exception on **top-level domain** for at least one of the **target**s specified in `arrayOfDomainStrings`.

The execution of this API and the use of the resulting permission (if granted) use the 'implicit' parameter, when the API is called, of the domain of the origin of the script (script-origin). If permission is granted, then the set of duplets (one per DOMstring):

```
[top-level-domain, DOMstring]
```

is added to the database of remembered grants.

A particular response to the API — like a DNT response header — is only valid immediately, and users' preferences may change.

A user agent MAY use an interactive method to ask the user about their preferences, so sites SHOULD NOT assume that the callback function will be called immediately.

6.4.2 API to cancel a site-specific exception

WebIDL

```
[NoInterfaceObject]
interface NavigatorDoNotTrack {
    void removeSiteSpecificTrackingException (sequence<DOMString> arrayOfDomainStrings);
};
```

6.4.2.1 Methods

`removeSiteSpecificTrackingException`

Ensures that the database of remembered grants no longer contains

```
[top-level-domain, DOMstring]
```

for all DOMstrings. This method never fails and there is no callback. After the call has been made, the indicated pairs are assured not to be in the database. The same matching as is used for determining which header to send is used to detect which entries (if any) to remove from the database.

> **Note**
>
> Note that establishing [site, *] and then requesting removal of [site, otherSite] simply leaves [site, *] in the database; the removal request has no effect and does **not** establish "grant an exception to everyone except otherSite".

Parameter	Type	Nullable	Optional	Description
arrayOfDomainStrings	`sequence<DOMString>`	✘	✘	

Return type: `void`

6.5 JavaScript API for web-wide exceptions

> **Issue**
>
> ISSUE-113: Should there be a JavaScript API to prompt for a Web-wide exception?
> **PROPOSAL**: In this section

6.5.1 API to request a web-wide exception

WebIDL

```
[NoInterfaceObject]
interface NavigatorDoNotTrack {
    void requestWebWideTrackingException (TrackingResponseCallback callback, optional  siteName, optional  explanationString, optional
};
```

6.5.1.1 Methods

`requestWebWideTrackingException`

If permission is granted, then the single duplet

`[ * , top-level-domain]`

is added to the database of remembered grants.

The parameters are as described above in the request for site-specific exceptions.

Parameter	Type	Nullable	Optional	Description
callback	`TrackingResponseCallback`	✘	✘	
siteName		✘	✔	
explanationString		✘	✔	
detailURI		✘	✔	

Return type: `void`

Users may wish to configure exceptions for a certain trusted tracker across all sites. This API requests the addition of a web-wide grant for a specific site, to the database.

6.5.2 API to cancel a web-wide exception

WebIDL

```
[NoInterfaceObject]
interface NavigatorDoNotTrack {
    void removeWebWideTrackingException ();
};
```

6.5.2.1 Methods

`removeWebWideTrackingException`

Ensures that the database of remembered grants no longer contains

`[ * , top-level-domain]`

This method never fails and there is no callback. After the call has been made, the indicated pair is assured not to be in the database. The same matching as is used for determining which header to send is used to detect which entry (if any) to remove from the database.

No parameters.
Return type: `void`

6.6 User interface guidelines

This section is non-normative.

User agents are free to implement exception management user interfaces as they see fit. Some agents might provide a prompt to the user at the time of the request. Some agents might allow users to configure this preference in advance. In either case, the user agent responds with the user's preference.

A user agent that chooses to implement a prompt to present tracking exception requests to the user might provide an interface like the following:

Example

```
Example News (web.exnews.com) would like to know
whether you permit tracking by a specific set of sites (click
here for their names).

Example News says:
  "These sites allow Example News to see how we're
  doing, and provide useful features of the Example News
  experience."     [More info]

[Allow Tracking]  [Deny Tracking Request]
```

In this example, the domains listed are those specified in `arrayOfDomainStrings`, the phrase "Example News" is from `siteName`, and the `explanationString` is displayed for the user with a "More info" link pointing to `detailURI`.

The user agent might then store that decision, and answer future requests based on this stored preference. A user agent might provide the user with an interface to explicitly remove (or add) user-granted exceptions.

Users might not configure their agents to have simple values for DNT, but use different browsing modes or other contextual information to decide on a DNT value. What algorithm a user agent employs to determine DNT values (or the lack thereof) is out of the scope of this specification.

In some user-agent implementations, decisions to grant exceptions may have been made in the past (and since forgotten) or may have been made by other users of the device. Thus, exceptions may not always represent the current preferences of the user. Some user agents might choose to provide ambient notice that user-opted tracking is ongoing, or easy access to view and control these preferences. Users may desire options to edit exceptions either at the time of tracking or in a separate user interface. This might allow the user to edit their preferences for a site they do not trust without visiting that site.

Issue

ISSUE-83: How do you opt out if already opted in?
PROPOSAL: In this section

Issue

ISSUE-129: Should a blanket exception of the type ["firstparty", "*"] be possible?
PROPOSAL: In this section

Issue

ISSUE-130: Should a global exception for a given third party on all sites be supported?
PROPOSAL: In this section

Issue

ISSUE-140: Do we need site-specific exceptions, i.e., concrete list of permitted third parties for a site?
PROPOSAL: In this section; yes, as some sites may have a mix of trusted/needed third parties, and others that either don't need to track, or aren't as trusted, or both.

6.7 Exceptions without a DNT header

Sites might wish to request exceptions even when a user arrives without a DNT header. Users might wish to grant affirmative permission to tracking on or by certain sites even without expressing general tracking preferences.

User agents MAY instantiate `NavigatorDoNotTrack.requestSiteSpecificTrackingException` even when `navigator.doNotTrack` is null. Sites SHOULD test for the existence of `requestSiteSpecificTrackingException` before calling the method. If an exception is granted in this context and the user-agent stores that preference, a user agent may send a DNT:0 header even if a tracking preference isn't expressed for other requests. Persisted preferences MAY also affect which header is transmitted if a user later chooses to express a tracking preference.

Note

Users might not configure their agents to have simple values for DNT, but use different browsing modes or other contextual information to decide on a DNT value. What algorithm a user agent employs to determine DNT values (or the lack thereof) is out of the scope of this specification.

6.8 Fingerprinting

By storing a client-side configurable state and providing functionality to learn about it later, this API might facilitate user fingerprinting and tracking. User agent developers ought to consider the possibility of fingerprinting during implementation and might consider rate-limiting requests or using other heuristics to mitigate fingerprinting risk. User agents SHOULD clear stored user-granted exceptions when the user chooses to clear cookies or other client-side state.

A. Acknowledgements

This specification consists of input from many discussions within and around the W3C Tracking Protection Working Group, along with written contributions from Nick Doty (W3C/MIT), Roy T. Fielding (Adobe), Tom Lowenthal (Mozilla), Jonathan Mayer (Stanford), Aleecia M. McDonald (Mozilla), Matthias Schunter (IBM), John Simpson (Consumer Watchdog), David Singer (Apple), Rigo Wenning (W3C/ERCIM), Shane Wiley (Yahoo!), and Andy Zeigler (Microsoft).

The DNT header field is based on the original *Do Not Track* submission by Jonathan Mayer (Stanford), Arvind Narayanan (Stanford), and Sid Stamm (Mozilla). The DOM API for `NavigatorDoNotTrack` is based on the *Web Tracking Protection* submission by Andy Zeigler, Adrian Bateman, and Eliot Graff (Microsoft). Many thanks to Robin Berjon for ReSpec.js.

B. References

B.1 Normative references

[ABNF]
D. Crocker and P. Overell. *Augmented BNF for Syntax Specifications: ABNF.* January 2008. Internet RFC 5234. URL: http://www.ietf.org/rfc/rfc5234.txt

[HTTP11]
R. Fielding; et al. *Hypertext Transfer Protocol - HTTP/1.1.* June 1999. Internet RFC 2616. URL: http://www.ietf.org/rfc/rfc2616.txt

[NAVIGATOR]
Ian Hickson, David Hyatt. *Navigator interface in HTML5.* 15 April 2011. Editors' draft. (Work in progress.) URL: http://dev.w3.org/html5/spec/timers.html#navigator

[RFC2119]
S. Bradner. *Key words for use in RFCs to Indicate Requirement Levels.* March 1997. Internet RFC 2119. URL: http://www.ietf.org/rfc/rfc2119.txt

[RFC4627]
D. Crockford. *The application/json Media Type for JavaScript Object Notation (JSON)* July 2006. Internet RFC 4627. URL: http://www.ietf.org/rfc/rfc4627.txt

[TRACKING-COMPLIANCE]
Justin Brookman; Sean Harvey; Erica Newland; Heather West. *Tracking Compliance and Scope.* 13 March 2012. W3C Working Draft. (Work in progress.) URL: http://www.w3.org/TR/2011/WD-tracking-compliance-20120313/

[WEBIDL]
Cameron McCormack. *Web IDL.* 27 September 2011. W3C Working Draft. (Work in progress.) URL: http://www.w3.org/TR/2011/WD-WebIDL-

20110927/

B.2 Informative references

[COOKIES]
Adam Barth. *HTTP State Management Mechanism*. April 2011. Internet Proposed Standard RFC 6265. URL: http://www.rfc-editor.org/rfc/rfc6265.txt

[KnowPrivacy]
Joshua Gomez; Travis Pinnick; Ashkan Soltani. *KnowPrivacy.* 1 June 2009. URL: http://www.knowprivacy.org/report/KnowPrivacy_Final_Report.pdf

[RFC5785]
Mark Nottingham; Eran Hammer-Lahav. *Defining Well-Known Uniform Resource Identifiers (URIs).* April 2010. Internet Proposed Standard RFC 5785. URL: http://www.rfc-editor.org/rfc/rfc5785.txt

[URI-TEMPLATE]
Joe Gregorio; Roy T. Fielding; Marc Hadley; Mark Nottingham; David Orchard. *URI Template.* March 2012. Internet RFC 6570. URL: http://www.rfc-editor.org/rfc/rfc6570.txt