Definition of context

As a follow-up question to the definition of “tracking” (http://www.w3.org/2011/tracking-protection/CfO_rationales/Tracking_decision.pdf) the group was asked if and what definition of “context” to include in the TPE specification.
The CfO was set up to solve and close the ISSUE-240: Do we need to define context? (http://www.w3.org/2011/tracking-protection/track/issues/240)

Based on the comments submitted, the Co-Chairs conclude that the Working Group has reached consensus to choose Option C as the least objectionable definition. This decision also accords with the group’s long understanding and previous resolution of ISSUE-10 (definition of parties) to circumscribe the meaning of a DNT signal to data collected across domains owned and operated by distinct entities.

The Call for Objections was open from February 12, 2014 to February 26, 2014. In total 14 members of the Working Group participated and presented arguments against or in favor of the four options. The full results of the questionnaire are public at https://www.w3.org/2002/09/wbs/49311/tpwg-context-240/results.

The Options

1. Option A: **Common controller and group identity**

   A context is a set of resources with a common data controller and a group identity that is easily discoverable by a user.

   Note that this definition of context is intended to represent a typical user's expectations regarding the boundaries of a commonly branded Web site (i.e., what makes it distinct from sites with a different group identity) independent of the technology, domain names, or parties operating that site via one or more origin servers.

2. Option B: **Data controller with common branding**

   A context is limited to the set of resources that share the same data controller, are covered by the same privacy policy, share a common branding, and whose host domains, other than that of the document origin, have been declared in the same-party property of the Tracking Resource.

   Non-normative Note:

   In case the same-party field is empty, then only the given site is considered to be the same context.

   In order for a definition of context to be granular enough to distinguish one context from another, a set of cumulative criteria is proposed. The purpose of this definition is to reflect the user expectations that data collected for a specified purpose by one of those resources is available to all other resources within the same context. Data must not be shared between different contexts. Respect for context and purpose limitation within a context are important core principles for any use of (personal) data within that context. Within any particular network interaction within a context, a user can expect that session states and other data (strictly) necessary to support the activity will be retained or shared.

3. Option C: **Resources operated by one party**

   A context is a set of resources that are controlled by the same party or parties.

   Note: This refers back to the working group's definition of party

4. Option D: **No definition**

   The notion of context is left undefined.

Explanatory considerations on the choice of definition

The decision was based on the substance of the objections against each option. The goal was to identify the option with the least substantiated objections. After evaluating the Working Group's inputs, we determined that Option D received substantial objections, especially as the definition of tracking (ISSUE-5) had previously been offered on the condition that context be defined. Therefore, it was first determined as the Working Group's decision to include a definition of context and not let the notion be undefined. Among the three text proposals for a definition, Options A, B, and C, the Options A and B received more substantial objections. Therefore, C is determined as the least objectionable definition for ISSUE-240. Additionally, Option C has the benefit that it is easily understandable, straightforward, and consistent with previous understanding and group decisions, linking the two essential concepts of cross-context tracking and party.

Objections against Option D:

Option D (no definition) received several substantial objections from the participants. The arguments can be summarized as such: as Option D leaves the most crucial element of the "tracking" definition undefined, it essentially leaves "tracking" itself undefined or only vaguely defined at best.

Mike O'Neill objected: "Because the definition of tracking relies on this notion, it must be defined."

Similarly, David Singer wrote: "This is only acceptable if the definition of tracking is used solely to inform the users roughly what it is that they have turned off. Even then, it is rather absurd for a definition to use a term that is both undefined and could be defined in such various ways. Fundamentally, this is pointing out we adopted a 'definition' which was anything but."

Roy Fielding also raised the same argument: "I object to this option because undefined terms just let others define them for you, often resulting in unnecessary legal actions."

Although a number of participants declined to support Option C for the reason of flexibility with regard to the potentially multiple compliance regimes defining "context" in different ways, the objections to Option D were more substantial.

The main reason to include a definition of "tracking" in the TPE was to enable the user to send a well-defined DNT signal being aware of the scope and coverage of what this signal conveys. Without a consistent, relatively precise definition of tracking, a "Do Not Track" signal would be less

impactful, and users would be less able to expect a consistent response from servers that track across multiple parties' sites. (See explanatory memo for the definition of tracking: http://www.w3.org/2011/tracking-protection/CfO_rationales/Tracking_decision.pdf)

By basing our definition of tracking on the concept of cross-context tracking, an essential part of the definition was left undefined. Therefore, ISSUE-240 was raised. Without the definition of context the initial reasoning to enable the user to send and understand a meaningful DNT signal would become inapplicable.

If on the other hand, the standard would allow any receiving server to define the notion of context (and thus, the definition of tracking), the user would send out a DNT signal without knowing in advance what his or her own preference would mean and what behavior would be encompassed. The initial goal of empowering the user to make a meaningful decision would be weakened by Option D. Consequently, the objections against Option D were more substantial than the objections to the other options.

Choice among Options A, B, and C:

The Options A, B, and C listed text proposals for a definition of "context". All three Options were were related to the concept of "data controller" or "party".

Option B appeared to be the most descriptive and strictest of the three Options: **"A context is limited to the set of resources that share the same data controller, are covered by the same privacy policy, share a common branding, and whose host domains, other than that of the document origin, have been declared in the same-party property of the Tracking Resource."**

This descriptiveness on requirements led to most objections raised against Option B. Moreover, this definition would be inconsistent with the group's previous decision on the definition of parry, which did not tie party identity to clear branding (http://www.w3.org/2011/tracking-protection/CfO_rationales/Party_decision.pdf).

Below listed are exemplary objections along these lines.

Chris Pedigo wrote: "I object to this option because it is not workable and would cause many companies to not implement this DNT standard. For example, many companies have different privacy policies for each region/country where they have a presence in order to reflect and honor that region or country's unique laws. Under this Option, a company's site in one country would be a different context than the company's nearly-identical site in another country. Finally, I am concerned that the non-normative text contains several phrases - "Data must not be shared", "respect for context and purpose limitation", and "data (strictly) necessary" - that would be inappropriate for the TPE and are better suited for a compliance regime."

Shane Wiley raised a similar objection: "it is a gross extension and overstatement of previously agreed upon positions with respect to party position."

David Singer objected: "This seems very restrictive, and tends too far into the domain of the compliance specifications. It's too complex to explain to a user, as well."

Rob Sherman raised several objections against Option B: "I'm concerned that Option B would cause parties to choose not to adopt the TPE, particularly on a global basis where adoption is especially important. Specifically, not all countries require privacy policies in the first instance, and

even where privacy policies are adopted or customary many companies will choose to adopt different privacy policies for services that operate in different countries, recognizing the varying compliance obligations that exist in different jurisdictions. Requiring a party to satisfy not one but all of these tests is unrealistic and inconsistent with consumer expectations — for example, if it is expected that a U.S.-based service and a non-U.S. service should be able to interoperate but have different privacy policies because of differing legal obligations. Likewise, the Working Group has already decided to reject common branding, which is unduly restrictive and does not cover all possible scenarios, as the appropriate standard for the Compliance document; it should not be reintroduced here.

Even if this issue were addressed, Option B is inferior to other options because it is based on the premise — as reflected in the reference to including all same-context domains in the same-party attribute — that context and party are coextensive or overlapping. This adds needless complexity to the TPE and creates confusion with two definitions of a similar concept that are not clearly distinguished, when in fact the definition of "party" can do all of the necessary work.

Option B also is defective because it is inconsistent with the Working Group's consensus decision in Section 2.3 of the TPE that some resources may be operated by multiple parties, a concept that is not incorporated into Option B.

Finally, the non-normative text included in Option B purports to prescribe compliance obligations (“reflect user expectations,” “data collected for a specified purpose ... is available to all other resources within the same context,” “[d]ata must not be shared between different contexts,” “[r]espect for context and purpose limitation ... are important core principles for any use of ... data,” etc.) that are inappropriate for inclusion in the TPE. These concepts, and the normative text they interpret, should be deferred to the appropriate compliance regime.”

We considered these arguments to be most decisive with regard to Option B. The requirement of the same privacy policy covering not only all services of a party but also one service in all available languages and countries seems overly restrictive. To accommodate different legislations it would not be unusual to have policies differing in regional details. To require to treat regional front ends of one service as different contexts could be a severe obstacle to implementation.

Rob Sherman’s and David Singer’s objection that Option B is unnecessarily complicated and hard to explain to users also holds true. While pointing in the same direction as the consensus definition of party, Option B adds additional requirements that lead to additional complexity that could impair comprehension of users and implementation of stakeholders. The competing Options A and C were both more closely tied with the party definition and therefore easier to understand and implement.

We recognize that this restrictiveness and descriptiveness of Option B could also be considered a benefit, as pointed out by Lee Tien: “Although I still am not sure what "data controller" means, this is the best of the options because it is more specific. The non-normative note also clarifies that the definition has the correct motivation -- to distinguish, rather than aggregate or blur, contexts. In particular, compared to C, a single party can operate different contexts and thus "track." This is a feature, not a bug”, Option A and Option C however were ultimately less objectionable to the participants and more aligned with our long standing consensus that first parties would be mostly excerpted from the specification.

Option A and C were very similar, with Option C explicitly referencing our party definition while Option A only party included the text from the party definition.

Option A reads: **“A context is a set of resources with a common data controller and a group identity that is easily discoverable by a user.”**

The party definition’s first sentence reads: “A party is a natural person, a legal entity, or a set of legal entities that share common owner(s), common controller(s), and a group identity that is easily discoverable by a user.”

Since the two Options A and C were similar in respect of content, Option A and Option C received similar or same objections. Several participants objected both Options as being too vague or too broad.

Exemplary are the following statements:

Walter van Holst objected both Options as being “vague”.

Rob van Eijk shares this concern about the vagueness of both Options: “I object, this definition does not help the definition of tracking. (1) On the contrary, it makes the meaning of tracking less clear. (2) It also lacks the granularity needed to distinguish different contexts between sets "of resources with a common data controller and a group identity that is easily discoverable by a user". E.g., a data controller with news magazine A (set of resources) and news magazine B (set of resources) could easily be one context under this definition.”

David Singer raised a similar argument: “this is too loose to be testable or usable””.

Mike O Neill wrote: “A data controller may have distinct privacy policies for data collected by servers for different sets of resources. A user may not agree that data collected under one policy can be linked to data collected under another, so the definition of context should reflect this.”

While both Options are certainly more vague than Option B, we do not share this concern as both Options clearly draw the line of “context” between different parties/data controllers. This concept is not inconsistent with existing legal regimes.

Both Options also seem to provide the desired descriptiveness with regard to the scope of the DNT signal sent by the user. As Roy Fielding pointed out: “The comments by others regarding vagueness are once again failing to understand the point. This notion of context needs to be understood by the user regardless of a site's compliance, since it is being sent without regard to any specific site.” The implementation and testability will then heavily rely on the accompanying compliance specification.

As both options received similar objections the decision between them was based on which Option appeared to be the least objectionable for the Working Group.

Several participants pointed out that while they object in general to include a definition of context in the TPE, in case a definition gets chosen anyway Option C would be the least objectionable to them. While the argument in favor of Option D needed to be dismissed because of the substantial objections against Option D, here the expressed preference was decisive. This preference for Option C among the three definitions was shared by Chris Pedigo, Chris Mejia, Brad Kulick, Shane Wiley, and Rob Sherman. All expressed that Option A and B were more objectionable to them than Option C, which was most aligned with our previous work and would tie two important concepts together, “tracking” and “party”.

Result

In conclusion, the ISSUE-240: Do we need to define context? (http://www.w3.org/2011/tracking-protection/track/issues/240) is hereby closed, and the following definition represents the Working Group's decision:

> **A context is a set of resources that are controlled by the same party or parties.**
>
> *Note: This refers back to the working group's definition of party*

Note: In the effort to finalize the TPE specification the two authors of Option C agreed to make a minor editorial change to the text of Option C to align it even more with the definition of party and to get rid of an unintentional ambiguity. The definition will get included in the TPE with the following wording:

> **A context is a set of resources that are controlled by the same party or jointly controlled by a set of parties.**