EFF POSITION PAPER: Unlinkability/auditability
W3C Workshop: Do Not Track and Beyond
26-27 November 2012

Defining "unlinkability" is a major issue for the Tracking Protection Working Group in crafting the TPE standard. On the W3C mailing list, Shane Wiley suggested a practical approach to handle some of the problems that have arisen in the TPWG discussions: roughly, that we develop a high-level definition as normative text, supplemented by non-normative text examples of aggressive approaches that clearly satisfy "unlinkability" (k-anonymity, URL filtering, super campaign structures, client-side storage, etc.).

EFF agrees that this may be a fruitful approach, but we believe it is critical to include non-normative text of approaches that would fail "unlinkability." The remainder of this position paper sets forth some of our current thinking.

Discussion

The current editor's draft defines "unlinkability" in two different ways. (The text set forth below corrects errors in that draft and thus is not quite verbatim.)

3.6.1 Option 1: Unlinkable Data

A party renders a dataset **unlinkable** when it
1. takes commercially reasonable steps to de-identify data such that there is confidence that it contains information which could not be linked to a specific user, user agent, or device in a production environment
2. publicly commits to retain and use the data in unlinkable fashion, and not to attempt to re-identify the data
3. contractually prohibits any third party that it transmits the unlinkable data to from attempting to re-identify the data. Parties should provide transparency to their delinking process (to the extent that it will not provide confidential details into security practices) so external experts and auditors can assess if the steps are reasonable given the particular data set.

3.6.2 Option 2: Unlinkable Data

A dataset is **unlinkable** when there is a high probability that it contains only information that could not be linked to a particular user, user agents, or device by a skilled analyst. A party renders a dataset unlinkable when either:
1. it publicly publishes information that is sufficiently detailed for a skilled analyst to evaluate the implementation, or
2. ensures that the dataset is at least 1024-unlinkable.

Discussion

We focus on the initial definition of "unlinkable," not the additional safeguards aimed

against re-identification attacks.

Option 1 is a variant of the Federal Trade Commission's (FTC) definition of "not reasonably linkable":

Data is not "reasonably linkable" to the extent that a company:
(1) takes reasonable measures to ensure that the data is de-identified;
(2) publicly commits not to try to re-identify the data; and
(3) contractually prohibits downstream recipients from trying to re-identify the data.

This definition turns on the meaning of "de-identified," which the FTC defined as: "the company must achieve a reasonable level of justified confidence that the data cannot reasonably be used to infer information about, or otherwise be linked to, a particular consumer, computer, or other device."

The FTC noted that this principle applies even when the data has not yet been linked to a particular consumer, computer or device, so long as it may reasonably become so linked. It also noted that if a company maintains and uses both identifiable data and data that has been de-identified, it should silo the two datasets separately. See generally Federal Trade Commission, Protecting Consumer Privacy in an Era of Rapid Change 20-22 (2012).

Option 1 is also similar to language used by the Digital Advertising Alliance.

A. De-Identification Process

Data has been De-Identified when an entity has taken reasonable steps to ensure that the data cannot reasonably be re-associated or connected to an individual or connected to or be associated with a particular computer or device.

An entity should take reasonable steps to protect the non-identifiable nature of data if it is distributed to non-Affiliates and obtain satisfactory written assurance that such entities will not attempt to reconstruct the data in a way such that an individual may be re-identified and will use or disclose the de-identified data only for uses as specified by the entity.

An entity should also take reasonable steps to ensure that any non-Affiliate that receives de-identified data will itself ensure that any further non-Affiliate entities to which such data is disclosed agree to restrictions and conditions set forth in this subsection V.A.

As with the FTC approach, the key language for present purposes is the actual definition of de-identification.

Significant questions have been raised about Option 1 on the mailing list. For instance, Shane Wiley suggested that "performing a one-way secret hash (salted hash) on identifiers (Cookie IDs, IP Addresses) and storing the resulting dataset in a logically/physically separate location from production data with strict access controls, policies, and employee education" would meet the definition set forth in Option 1. His

stated goal was to "find the middle-ground between complete destruction of data and an unlinkable state that still allows for longitudinal consistency for analytical purposes BUT CANNOT be linked back to a production system such that the data could be used to modify a single user's experience."

Prof. Felten, however, argued: "hashing IP addresses (with or without salting) does not render them unlinkable. After hashing, it's easy to recovery the original IP address. The story is similar for other types of unique identifiers--there are ways to get to unlinkability, but hashing by itself won't be enough."

Option 2, which derives from the EFF/Stanford/Mozilla proposal, differs mainly from Option 1 by requiring a "high probability" rather than "reasonableness," defined either objectively (1024-unlinkable) or by expert analysis.

This approach is similar to the treatment of health data under the HIPAA de-identification rule. HIPAA establishes a high standard that patient information is "de-identified" only when "there is no reasonable basis to believe that the information can be used to identify an individual." 45 C.F.R. § 164.514 (a). The current HIPAA rule then provides two "safe harbors": data is de-identified if one follows a prescribed approach of removing identifiers, or if certified as de-identified by a statistician.

The first safe harbor involves the removal of eighteen specified patient identifiers, including but not limited to, patient name, location (other than state or 3-digit ZIP codes with populations greater than 20,000), email address, telephone number, Social Security Number, and the like. 45 C.F.R. § 164.514(b)(2)(i). Significantly, the eighteenth identifier that must be removed is "any other unique identifying number, characteristic, or code." Prof. Sweeney's research suggests that under this safe harbor, about 0.04% of the population can be re-identified.
http://dataprivacylab.org/projects/identifiability/pharma1.pdf

The second safe harbor requires a formal determination by a qualified statistician who, applying statistical and scientific principles and methods for rendering information not individually identifiable, determines that the risk is very small that the information could be used, alone or in combination with other reasonably available information, by an anticipated recipient to identify an individual who is a subject of the information. The statistician must reach a conclusion that the risks of re-identification are "very small" in order for the patient information to be properly "de-identified." Unfortunately, "very small" is not defined. An obvious suggestion here is that "very small" be taken to mean identifiability of 0.04%, analogous to Option 2's 1024-unlinkable rubric.

Auditability

A further weakness of the "statistician" safe harbor under HIPAA is that there is no oversight of the statistician's "very small" risk determination. Even for good actors, opacity of the de-identification techniques and assumptions about background data can make it hard to know whether a risk that was "very small" will become higher as data and

re-identification techniques improve.

Option 1 and Option 2 each attempt to address this problem. Under Option 1, parties should provide transparency to their delinking process (to the extent that it will not provide confidential details into security practices) so external experts and auditors can assess if the steps are reasonable given the particular data set.

Under Option 2, the alternative to 1024-unlinkability is that the entity claiming de-identification publicly publish information that is sufficiently detailed for a skilled analyst to evaluate the implementation.

We believe that some form of these requirements is critical to assuring that data remains unlinkable under the standard.

Conclusion

It may be possible to define "unlinkability" along the lines suggested by Shane Wiley: a high-level definition in normative text, with non-normative text setting forth examples of acceptable or even aggressive implementations. In order to set a floor, however, we believe that the non-normative text must also include examples of technical methods that are not acceptable. In addition, the standard must contain a meaningful process for auditing/verification.