Digging into the Privacy Sandbox

- Summary
- The current state of privacy on the web {: #background }
- Introducing the Privacy Sandbox {: #introduction }
  - The Privacy Sandbox proposals {: #proposals }
- Use cases and goals
  - Measure conversion
  - Select ads
  - Combat fingerprinting
  - IP address security
  - Combat spam, fraud and denial-of-service attacks
  - Enable domains to belong to the same first party
- Find out more
  - Privacy Sandbox proposal explainers {: #explainers }
  - The Privacy Sandbox
  - Discussion and participation
  - Use cases, policies, and requirements {: #use-cases-policies-requirements }
- Appendix: Glossary of terms used in the proposal explainers {: #glossary }
  - Click-through rate (CTR) {: #glossary-ctr }
  - Click-through-conversion (CTC) {: #glossary-ctc }
  - Conversion
  - Differential privacy
  - Domain
  - eTLD, eTLD+1 {: #glossary-etld }
  - Entropy
  - Fingerprinting {: #glossary-fingerprinting }
  - Fingerprinting surface {: #glossary-fingerprinting-surface }

- First-party {: #glossary-first-party }
- Impression {: #glossary-impression }
- k-anonymity
- Nonce
- Origin
- Passive surface {: #glossary-passive-surface }
- Publisher
- Reach
- Remarketing
- Site
- Surface
- Third-party {: #glossary-third-party }
- Top-level domain (TLD) {: #glossary-tld }
- .well-known

---

https://www.youtube.com/watch?v=WnCKlNE52tc

Summary

- This post outlines APIs and concepts from the Privacy Sandbox proposals.
- The proposal authors are inviting feedback from the community, particularly from those in the advertising space (publishers, advertisers, and ad tech companies), to suggest missing use cases and share information about how to support your business use cases.
- You can comment on the proposals by filing issues on the repositories linked to below.
- There's a glossary for the proposals at the end of this post.

The current state of privacy on the web {: #background }

Websites use services from other companies to provide analytics, serve video and do lots of other useful stuff. Composability is one of the web's superpowers. Most notably, ads are included in web pages via third-party JavaScript and iframes. Ad views, clicks and conversions are tracked via third-

party cookies and scripts.

However, when you visit a website you may not be aware of the third parties involved and what they're doing with your data. Even publishers and web developers may not understand the entire third-party supply chain.

Ad selection, conversion measurement, and other use cases currently rely on establishing stable cross-site user identity. Historically this has been done with third-party cookies, but browsers have begun to restrict access to these cookies. There has also been an increase in the use of other mechanisms for cross-site user tracking, such as covert browser storage, device fingerprinting, and requests for personal information such as email addresses.

This is a dilemma for the web. How can legitimate third-party use cases be supported without enabling users to be tracked across sites?

In particular, how can websites fund content by enabling third parties to show ads and measure ad performance—but not allow individual users to be profiled? How can advertisers and site owners evaluate a user's authenticity without resorting to dark patterns such as device fingerprinting?

The way things work at the moment can be problematic for the entire web ecosystem, not just users. For publishers and advertisers, tracking identity and using a variety of non-standard third-party solutions can add to technical debt, code complexity, and data risk. Users, developers, publishers, and advertisers should be confident that the web is protecting user privacy choices.

Advertising is a core web business model for the internet, but advertising has to work for everyone. Which brings us to the Privacy Sandbox's mission: to create a thriving web ecosystem that is respectful of users and private by default.

Introducing the Privacy Sandbox {: #introduction }

The Privacy Sandbox introduces a set of privacy-preserving APIs to support business models that fund the open web in the absence of tracking mechanisms like third-party cookies.

The Privacy Sandbox APIs require web browsers to take on a new role. Rather than working with limited tools and protections, the APIs enable the user's browser to act on the user's behalf—locally, on their device—to protect the user's identifying information as they navigate the web. The APIs enable use cases such as ad selection and conversion measurement, without revealing individual private and personal information. In engineering terms a sandbox is a protected environment; a key principle of the Privacy Sandbox is that a user's personal information should be protected and not shared in a way that lets the user be identified across sites.

This is a shift in direction for browsers. The Privacy Sandbox's vision of the future has browsers providing specific tools to satisfy specific use cases, while preserving user privacy. A Potential Privacy Model for the Web sets out core principles behind the APIs:

- To establish the range of web activity across which the user's browser can let websites treat a person as having a single identity.
- To identify the ways in which information can move across identity boundaries without compromising that separation.

The Privacy Sandbox proposals {: #proposals }

In order to successfully transition away from third-party cookies the Privacy Sandbox initiative needs your support. The proposal explainers need feedback from developers as well as publishers, advertisers, and ad technology companies, to suggest missing use cases and share information about how to accomplish their goals in a privacy-safe way.

You can comment on the proposal explainers by filing issues against each repository:

- Privacy Model for the Web
  Establish the range of web activity across which the user's browser can let websites treat a person as having a single identity. Identify the ways in which information can move across identity boundaries without compromising that separation.
- Privacy Budget
  Limit the total amount of potentially identifiable data that sites can access. Update APIs to reduce the amount of potentially identifiable data revealed. Make access to potentially identifiable data measurable.
- Gnatcatcher
  Limit the ability to identify individual users by accessing their IP address.
- Trust Token API
  Enable an origin that trusts a user to issue them with cryptographic tokens which are stored by the user's browser so they can be used in other contexts to evaluate the user's authenticity.
- First-Party Sets
  Allow related domain names owned by the same entity to declare themselves as belonging to the same first party.
- Aggregated Reporting
  Provide privacy preserving mechanisms to support a variety of use cases such as view-through-conversion, brand, lift, and reach measurement.

- Attribution Reporting
  Provide privacy preserving measurement of clicks and views with event-level and aggregate reports.
- Topics API
  Enable interest-based advertising, without having to resort to tracking the sites a user visits. The design of the API was informed by community feedback from our earlier FLoC trials, and supersedes the FLoC proposal.
- FLEDGE
  Provides a solution for remarketing use cases, designed so it cannot be used by third parties to track user browsing behaviour across sites.

You can dive into the API proposal explainers right away, and over the coming months we'll be publishing posts about each proposal individually.

We'll also be adding to our playlist of five-minute videos that give a simple explanation of each API.

https://www.youtube.com/embed/videoseries?list=PLNYkxOF6rcICntazGfSVKSj5EwuR9w5Nv

Use cases and goals

Measure conversion

Goal: Enable advertisers to measure ad performance.

The Attribution Reporting API enables the measurement of two events that are linked together:

1. An event on a publisher's website, such as a user viewing or clicking an ad.
2. A subsequent conversion on an advertiser site.

This API supports **click-through** and **view-through** measurement.

Other features in this API include cross-device attribution reporting and app-to-web attribution reporting.

The API also offers two types of **attribution reports**:

- **Event-level reports** associate a particular ad click or view (on the ad side) with data on the conversion side. To preserve user privacy, by preventing the joining of user identity across sites, conversion-side data is very limited, and the data is 'noised' (meaning that for a small percentage of cases, random data is sent). As an extra privacy protection, reports are not sent immediately.

- **Aggregate reports** are not tied to a specific event on the ad side. These reports provide richer, higher-fidelity conversion data than event-level reports. A combination of privacy techniques across cryptography, distribution of trust, and differential privacy help reduce the risk of identity joining across sites.

Both report types can be used simultaneously: they're complementary.

Introduction to Attribution Reporting explains more about the status of these features and how to try this API.

Select ads

Goal: Enable advertisers to display ads relevant to users.

Relevant ads are more favorable to users and more profitable for publishers (the people running ad-supported websites). Third party ad selection tools make ad space more valuable to advertisers (the people who purchase ad space on websites) which in turn increases revenue for ad-supported websites and enables content to get created and published.

There are many ways to make ads relevant to the user, including the following:

- **First-party-data**: Show ads relevant to topics a person has told a website they have an interest in, or content a person has looked at previously on the current website.
- **Contextual**: Choose where to display ads based on site content. For example, 'Put this ad next to articles about knitting.'
- **Remarketing**: Advertise to people who've already visited your site, while they are not on your site. For example, 'Show this ad for discount wool to people who visited your store and left knitting items in their shopping cart—while they're visiting craft sites.'
- **Interest-based**: Select ads based on a user's browsing history. For example, 'Show this ad to users whose browsing behaviour indicates they might be interested in knitting'.

First-party-data and contextual ad selection can be achieved without knowing anything about the user other than their activity within a site. These techniques don't require cross-site tracking.

Remarketing is usually done by using cookies or some other way to recognize people across websites: adding users to lists and then selecting specific ads to show them.

Interest-based ad selection currently uses cookies to track user behaviour across as many sites as possible. Many people are concerned about the privacy implications of ad selection. The Privacy Sandbox proposes two alternatives, for remarketing and for interest-based selection:

- FLEDGE: for **remarketing use cases**.
  The API is designed so it cannot be used by third parties to track user browsing behaviour: the user's browser, not the advertiser or ad tech platform, stores advertiser-defined

interest groups that the user's browser is associated with. The user's browser combines interest group data with ad buyer/seller data and business logic to conduct an "auction" locally on the user's device to select an ad, rather than sharing data with a third party.

- Topics API: for **interest-based audiences**.
  Enable interest-based advertising, without having to resort to tracking the sites a user visits. The API proposes using machine learning to infer topics from hostnames, and a JavaScript API that returns coarse-grained topics a user might currently be interested in, based on the hostnames of sites recently visited.

Combat fingerprinting

Goal: Reduce the amount of potentially identifiable data revealed by APIs and make access to potentially identifiable data controllable by users, and measurable.

Browsers have taken steps to deprecate third-party cookies, but techniques to identify and track the behaviour of individual users, known as fingerprinting, have continued to evolve. Fingerprinting uses mechanisms that users aren't aware of and can't control.

- The Privacy Budget proposal aims to limit the potential for fingerprinting by identifying how much fingerprint data is exposed by JavaScript APIs or other 'surfaces' (such as HTTP request headers) and setting a limit on how much of this data can be accessed.
- Fingerprinting surfaces such as the User-Agent header will be reduced in scope, and the data made available by alternative mechanisms such as Client Hints will be subject to Privacy Budget limits. Other surfaces, such as the device orientation and battery-level APIs, will be updated to keep the information exposed to a minimum.

IP address security

Goal: Control access to IP addresses to reduce covert fingerprinting, and allow sites to opt out of seeing IP addresses in order to not consume privacy budget.

A user's IP address is the public 'address' of their computer on the internet, which in most cases is dynamically assigned by the network through which they connect to the internet. However, even dynamic IP addresses may remain stable over a significant period of time. Not surprisingly, this means that IP addresses are a significant source of fingerprint data.

The Gnatcatcher proposal is an attempt to provide a privacy-preserving approach that avoids consuming privacy budget, while ensuring that sites requiring access to IP addresses for legitimate purposes such as abuse prevention can do so, subject to certification and auditing.

There are two parts to the proposal:

- **Willful IP Blindness** provides a way for websites to let browsers know they are not connecting IP addresses with users.
- **Near-path NAT** allows groups of users to send their traffic through the same privatizing server, effectively hiding their IP addresses from a site host.

Combat spam, fraud and denial-of-service attacks

Goal: Verify user authenticity without fingerprinting.

Anti-fraud protection is crucial for keeping users safe, and to ensure that advertisers and site owners can get accurate ad performance measurements. Advertisers and site owners must be able to distinguish between malicious bots and authentic users. If advertisers can't reliably tell which ad clicks are from real humans, they spend less, so site publishers get less revenue. Many third party services currently use techniques such as device fingerprinting to combat fraud.

Unfortunately, the techniques used to identify legitimate users and block spammers, fraudsters, and bots work in ways similar to fingerprinting techniques that damage privacy.

- The Trust Tokens API proposes an alternative approach, allowing authenticity established for a user in one context, such as a social media site, to be conveyed to another context, such as an ad running on a news site—without identifying the user or linking the two identities.

Enable domains to belong to the same first party

Goal: Enable entities to declare that related domain names are owned by the same first party.

Many organizations own sites across multiple domains. This can become a problem if restrictions are imposed on tracking user identity across sites that are seen as 'third-party' but actually belong to the same organization.

- First Party Sets aims to make the web's concept of first and third parties more closely aligned with the real world's by enabling multiple domains to declare themselves as belonging to the same first party.

Find out more

Privacy Sandbox proposal explainers {: #explainers }

The Privacy Sandbox initiative needs your support. The API proposal explainers need feedback, in particular to suggest missing use cases and more-private ways to accomplish their goals.

- Privacy Budget
- Trust Token API
- Gnatcatcher
- Aggregated Reporting API
- Conversion measurement
- Topics API
- FLEDGE

A Potential Privacy Model for the Web sets out the core principles underlying the APIs.

The Privacy Sandbox

- Technical overview: The Privacy Sandbox
- Non-technical overview: privacysandbox.com
- Project principles: Building a more private web
- The future of third-party cookies

Discussion and participation

- How to participate in the Privacy Sandbox initiative
- Progress update on the Privacy Sandbox initiative
- Privacy Sandbox Developer Support: join discussions or ask questions.

Use cases, policies, and requirements {: #use-cases-policies-requirements }

- Advertising Use Cases
- Mozilla anti-tracking policy
- WebKit tracking prevention policy
- Privacy Preserving Ad Click Attribution For the Web
- Brave, Fingerprinting, and Privacy Budgets

Appendix: Glossary of terms used in the proposal explainers {: #glossary }

Click-through rate (CTR) {: #glossary-ctr }

The ratio of users who click on an ad, having seen it. (See also impression.)

Click-through-conversion (CTC) {: #glossary-ctc }

A conversion attributed to an ad that was 'clicked'.

Conversion

The completion of an action on an advertiser's website by a user who has previously interacted with an ad from that advertiser. For example, purchase of a product or sign-up for a newsletter after clicking an ad that links to the advertiser's site.

Differential privacy

Share information about a dataset to reveal patterns of behaviour without revealing private information about individuals or whether they belong to the dataset.

Domain

See Top-Level Domain and eTLD.

eTLD, eTLD+1 {: #glossary-etld }

'Effective' top level domains are defined by the Public Suffix List. For example:

```
co.uk
appspot.com
glitch.me
```

Effective TLDs are what enable foo.appspot.com to be a different site from bar.appspot.com. The effective top-level domain (**eTLD**) in this case is appspot.com, and the whole **site** name (foo.appspot.com, bar.appspot.com) is known as the **eTLD+1**.

See also Top-Level Domain.

Entropy

A measure of how much an item of data reveals individual identity.

Data entropy is measured in bits. The more that data reveals identity, the higher its entropy value.

Data can be combined to identify an individual, but it can be difficult to work out whether new data adds to entropy. For example, knowing a person is from Australia doesn't reduce entropy if you already know the person is from Kangaroo Island.

Fingerprinting {: #glossary-fingerprinting }

Techniques to identify and track the behaviour of individual users. Fingerprinting uses mechanisms that users aren't aware of and can't control. Sites such as Panopticlick and amiunique.org show how fingerprint data can be combined to identify you as an individual.

Fingerprinting surface {: #glossary-fingerprinting-surface }

Something that can be used (probably in combination with other surfaces) to identify a particular user or device. For example, the `navigator.userAgent()` JavaScript method and the `User-Agent` HTTP request header provide access to a fingerprinting surface (the user agent string).

First-party {: #glossary-first-party }

Resources from the site you're visiting. For example, the page you're reading is on the site web.dev and includes resources from that site. See also Third-party.

Impression {: #glossary-impression }

View of an ad. (See also click-through rate.)

k-anonymity

A measure of anonymity within a data set. If you have *k* anonymity, you can't be distinguished from *k-1* other individuals in the data set. In other words, *k* individuals have the same information (including you).

Nonce

Arbitrary number used once only in cryptographic communication.

Origin

The origin of a request, including the server name but no path information. For example: `https://web.dev`.

Passive surface {: #glossary-passive-surface }

Some fingerprinting surfaces, such as user agent strings, IP addresses and accept-language headers, are available to every website whether the site asks for them or not. That means passive surfaces can easily consume a site's privacy budget.

The Privacy Sandbox initiative proposes replacing passive surfaces with active ways to get specific information, for example using Client Hints a single time to get the user's language rather than having an accept-language header for every response to every server.

Publisher

The Privacy Sandbox proposal explainers are mostly about ads, so the kinds of publishers referred to are ones that put ads on their websites.

Reach

The total number of people who see an ad.

Remarketing

Advertising to people who've already visited your site. For example, an online store could show ads for a toy sale to people who previously viewed toys on their site.

Site

See Top-Level Domain and eTLD.

Surface

See Fingerprinting surface and Passive surface.

Third-party {: #glossary-third-party }

Resources served from a domain that's different from the website you're visiting. For example, a website foo.com might use analytics code from google-analytics.com (via JavaScript), fonts from use.typekit.net (via a link element) and a video from vimeo.com (in an iframe). See also First-party.

Top-level domain (TLD) {: #glossary-tld }

Top-level domains such as .com and .org are listed in the Root Zone Database.

Note that some 'sites' are actually just subdomains. For example, translate.google.com and maps.google.com are just subdomains of google.com (which is the eTLD + 1).

.well-known

It can be useful to access policy or other information about a host *before* making a request. For example, robots.txt tells web crawlers which pages to visit and which pages to ignore. IETF RFC8615 outlines a standardized way to make site-wide metadata accessible in standard locations in a /.well-known/ subdirectory. You can see a list of these at iana.org/assignments/well-known-uris/well-known-uris.xhtml.

Thanks to all those who helped with writing and reviewing this post.

Photo by Pierre Bamin on Unsplash.

□□□ #3
□ □□□ □□□ 2026-10-31 21:50:01 UTC
□ cocobook-sync □□□ 2026-11-04 17:37:03 UTC