API Snapshots: Deep Dive

Privacy Sandbox

FedCM

Federated Credential Management

Facilitate seamless sign-in while safeguarding privacy

What's the privacy challenge?

Federated credentials allow people to sign in to various sites using their existing logins from Identity Providers, like social logins. This process enhances security and convenience for people while streamlining login management for businesses.

Currently, when someone visits a website with multiple login options, all of those providers may be notified about the visit, even if the user doesn't use their service to sign in. And some companies use this information to build out user profiles for advertising purposes.

Do things differently with FedCM

Make "Sign in with..." login solutions more private by shielding people's information to prevent passive tracking

User experience: FedCM helps streamline login solutions, leading to a smoother browsing experience for users.

Privacy-preserving data exchange: The login experience is handled by the browser, meaning that no information is shared with the site until a user wants it to be. Identity Providers also remain unaware of the user's online activity unless the user logs in with the service.

How it works

Step 1: Someone wants to log in to a site

Someone visits a site which requires them to log in or verify their identity for certain actions, like making a purchase or accessing premium content.

Step 2: FedCM is initiated

The site (known as the 'Relying Party') uses FedCM to allow people to easily sign in with an existing account from a trusted Identity Provider. The Relying Party chooses which Identity Providers to support on their site.

Step 3: Identity Provider is selected

The user selects their preferred Identity Provider from a dialog box in the browser. The site can't see the Identity Provider, and the Identity Providers can't see the site. The user's information stays with the browser until they log in with a service.

Step 4: Information is exchanged

Once the user initiates a login, the browser securely transmits the user's account selection to the Identity Provider, which verifies user credentials and generates a token. The token is validated, and the browser relays it to the Relying Party. At this point, FedCM steps back, and the Relying Party and Identity Provider can share information directly with each other. The login credentials are not shared, but pertinent information like email address or home address can be shared for a more seamless site experience.

Step 5: The user is logged in and can continue their task on site

The user is now logged in to the site. Only the Identity Provider selected for login can communicate with the Relying Party, reducing the amount of information shared without the user's knowledge.

Learn more at goo.gle/fedcm

Privacy Sandbox

Related Website Sets

Facilitate seamless experiences across related sites while safeguarding privacy

What's the privacy challenge?

Current online tracking practices allow companies to collect personal data across sites, often without the person's knowledge, to build detailed user profiles.

Do things differently with Related Website Sets

Enable limited data sharing within small groups of affiliated sites that meet specific criteria, preserving site experiences while preventing tracking and data leakage

User experience: Within a limited network of affiliated sites, first-party data sharing can be carefully managed to preserve essential site functionality and user experiences.

Privacy-preserving data exchange: Related Website Sets allow data sharing only within a well-defined set of affiliated websites, determined by explicit declarations and a user's interaction with the websites. This limits cross-site tracking and reduces data leakage.

How it works

Step 1: Someone browses the web

A person visits a site owned by Organization A. They take some type of action, like clicking on a widget or embed, which acts as a nod to the browser to check whether the site is part of a Related Website Set.

Step 2: Related Website Set is identified

The browser checks if the site belongs to a Related Website Set using locally stored data. Sets can include up to six affiliated sites: one primary site and five related sites such as news-organizationA.com and shopping-organizationA.com. Sites in a set can have unlimited geographical top-level domains (like .com and .co.uk) and unlimited service domains (like cdn-organizationA.com) that support site functionality. *Note: a site cannot be in multiple sets. This ensures that data cannot be joined across multiple sets and limits tracking potential.*

Step 3: Data sharing boundaries are established

For sites within a Related Website Set, a declaration is made specifying the legitimate relationships between them, including recognizable connections or service domain functionality. The browser then evaluates this declaration to establish and enforce data sharing boundaries within the set.

Step 4: Site capabilities are enabled

With data sharing boundaries established, sites within the set can leverage shared information to enhance user experience. This includes streamlining logins across related platforms, providing personalized experiences through shared settings and preferences, and coordinating actions like content loading.

Learn more at goo.gle/rws

Privacy Sandbox

CHIPS

Cookies Having Independent Partitioned State

Facilitate seamless site experiences while safeguarding privacy

What's the privacy challenge?

Many website features rely on embedded third-party tools. These tools often use cookies to store site visitor information such as session IDs, preferences, and analytics data. However, this data can be accessed by a third-party across any website where their tools are embedded, enabling widespread and potentially unanticipated cross-site tracking.

Do things differently with CHIPS

Maintain embedded site functionality by allowing access only to specific, compartmentalized information in a way that prevents tracking

Partitioned storage: CHIPS works by partitioning cookie storage based on the website's top-level domain. This prevents embedded third-party tools from accessing cookies across different sites, effectively limiting cross-site tracking.

Developer flexibility: Developers have granular control over cookie partitioning with CHIPS. They can choose to partition cookies individually, linking specific cookies to the top-level site when it benefits the user experience.

How it works

Step 1: Someone visits a site, prompting cookie creation and storage

Someone visit a website, where an embedded third-party service sets a cookie on their device. This cookie stores information about the person and their interactions with the site—like shopping cart items or language preferences—to improve their experience.

Step 2: CHIPS isolates and restricts access to the cookie

CHIPS isolates the cookie within a secure, site-specific container, like an individual cookie jar with the lid firmly sealed. This cookie jar can only be opened by the third-party provider on the same top-level site where the cookie was created, and only that third-party can read, modify, or delete the cookie.

The cookie is kept entirely separate from data used by other third-party features, like embedded content, and prevents the third-party who set the cookie from reading it on other sites the user may visit.

Step 3: Privacy is protected while maintaining relevant web experiences

This isolation prevents unintended sharing and tracking across different websites, safeguarding people's information while preserving relevant experiences on individual sites.

Learn more at goo.gle/chips

Privacy Sandbox

Topics

Reach relevant audiences powered by anonymized insights

What's the privacy challenge?

While third-party cookies help enable relevant advertising, they can also expose people's online behavior and personal information to a number of companies throughout the advertising process.

Do things differently with Topics

Reach interest-based audiences informed by users' recently visited sites, without sharing their specific browsing history across the web

Relevant ads: Advertisers can still deliver relevant ads, but relevance is based on broader interests rather than personal information.

On-device processing: Topics are determined directly on the user's device based on their browsing history, ensuring sensitive individual-level data stays on the device and is not shared with external parties.

How it works

Step 1: Sites are categorized according to topics
The browser uses a built-in taxonomy model to analyze website hostnames and categorize websites into topics like "sports," "travel," or "finance."

Step 2: The browser uses Topics to gauge people's current interests
When someone browses the internet, the browser analyzes the sites they visit within a seven-day window to infer five topics representing their top interests. These topics are updated each week based on ongoing browsing activity.

Step 3: The ad platform leverages Topics to serve relevant ads
To help an ad platform determine which ad to serve to a site visitor, the browser shares a single topic from each week, across a maximum of three weeks. To further enhance privacy, five percent of the topics are random.

This enables the ad platform to serve a relevant ad to the user, without revealing their specific browsing history or enabling the creation of overly detailed user profiles.

Learn more at goo.gle/topics

Privacy Sandbox

Protected Audience API

Re-engage relevant audiences powered by anonymized insights

What's the privacy challenge?

Third-party cookies enable advertisers to target audiences with relevant ads, but people's browsing behavior and personal information can be exposed to numerous companies throughout the process.

Do things differently with Protected Audience API

Create custom audiences for use cases like remarketing, without exposing personal information

Relevant ads: Advertisers and publishers can deliver relevant ads to past site visitors without relying on individual tracking.

On-device auctions: Auctions run on the user's device, providing an efficient bidding process while eliminating exposure of sensitive information.

How it works

Step 1: A site visitor is assigned to interest groups

To classify their site visitors, brands, or their ad techs can set criteria for interest groups, using factors like time spent on pages or products viewed. The browser then analyzes on-site behavior and anonymously assigns site visitors to relevant groups.

For example, if someone visits a fashion brand's website and browses their shoe collection, they could be added to an interest group called "in market for shoes". Importantly, the user's individual information remains on their device and is not shared back to the brand or any other server.

Step 2: The brand bids to show ads to specific interest groups, not specific people

Later, when this person visits a news website with an available ad slot, the site initiates a Protected Audience ad auction on-device. The advertiser can then bid on the chance to reach the interest group, rather than an individual identifiable person. This allows the fashion brand to re-engage the user with an ad for shoes, all without the news site or the advertiser knowing the user's identity or their specific browsing history.

Step 3: A relevant ad is selected and delivered

The winning ad is displayed on the user's device, and the final outcome of the bid is reported back to the SSP and the winning DSP.

This process helps safeguard users' personal information throughout the advertising process.

Learn more at goo.gle/pa

Privacy Sandbox

Attribution Reporting API

Enable effective ad measurement with privacy-preserving technologies

What's the privacy challenge?

Today, third-party cookies are widely used to track people across websites and attribute conversions to specific ads or campaigns for measurement. However, this practice can passively expose people's data and cross-site behavior to a large number of companies.

Do things differently with Attribution Reporting

Measure which marketing efforts are most effective by matching ad interactions with conversions, without identifying people at an individual level

Effective insights: Advertisers and publishers gain insights into campaign effectiveness without relying on individual-level information. Ad tech companies use the insights to power bidding models and drive key outcomes.

Privacy techniques: Attribution Reporting leverages privacy techniques like aggregation, noise addition, and on-device processing to deliver accurate attribution data while protecting user privacy.

How it works

Step 1: Someone sees and interacts with an online ad

Someone browses the web, visits a site with ads, and views or clicks on an ad. Ad tech on the website use the Attribution Reporting API to record that interaction in the browser.

Step 2: The person takes an action on the website

The person then takes a desired action on the advertiser's website, like making a purchase, and this conversion is recorded in the browser.

Step 3: Reporting is generated

Two types of reports can be generated to help advertisers and ad tech providers understand performance. Privacy-preserving techniques like aggregation, noise, and encryption are used to protect individual identity.

Event-level Reports
For ad specific insights

Provide visibility into individual ad events that drive conversions, including: ad type (e.g., banner ad), engagement type (e.g., click, view), and conversion type (e.g., purchase, registration)

Limited ad and conversion metadata sent on a schedule of report windows; flexible down to one hour or up to 30 days

Summary Reports
For aggregate level insights

Provide detailed and flexible aggregated performance reporting, including ROI analysis, campaign performance, and total purchase values by product category

Detailed conversion insights sent instantly or with random delay between zero and 10 minutes

Learn more at goo.gle/ar