



Privacy Sandbox Progress Report

Q3 Reporting Period - July to September 2023

Prepared for the CMA, 24 October 2023

Overview

Google has prepared this quarterly report as part of its Commitments to the Competition and Markets Authority ('CMA') under paragraphs 12, 17(c)(ii) and 32(a). This report covers Google's progress on the Privacy Sandbox proposals; updated timing expectations; substantive explanations of how Google has taken into account observations made by third parties; and a summary of interactions between Google and the CMA, including feedback from the CMA and Google's approach to addressing the feedback.

Progress of Privacy Sandbox Proposals

Google has been keeping the CMA updated on progress with the Privacy Sandbox proposals in its regular Status Meetings scheduled in accordance with paragraph 17(b) of the Commitments. Additionally, the team maintains the [Privacy Sandbox developer documentation](#) with specific pages for each API, an overall [status page](#), along with continued updates on core project processes such as [Chrome-facilitated testing](#) and [preparing for third-party cookie deprecation](#). Key updates are shared under [the "Privacy" tag on the developer blog](#) along with targeted updates shared to the individual developer mailing lists.

Updated Timing Expectations

Google's latest expectations for the timing of the Privacy Sandbox proposals are set out in the [Privacy Sandbox Timeline](#).¹ The summary below includes all Q3 2023 updates, covering the period from July 11 to September 30, 2023.

¹ According to Annex 1 of the Commitments, if the development of an API is discontinued and/or alternative APIs developed, such changes will be reported and reflected in Google's public updates, as provided for in paragraph 11 of the Commitments. Under paragraph 17(a) of the Commitments, Google is required to proactively inform the CMA of changes to the Privacy Sandbox that are material and without delay seek to resolve concerns raised and address comments made by the CMA with a view to achieving the Purpose of the Commitments.

Privacy Sandbox Q3 2023 Timeline Updates	
July Timeline Updates	• No updates
August Timeline Updates	• Added “OT closed” for: ○ Topics API ○ Protected Audience API ○ Attribution Reporting API ○ Shared Storage API ○ Fenced Frames API • Updated tooltip copy: ○ Topics API: The origin trial for Topics API ran from Chrome 101 to 115. ○ Protected Audience API: The origin trial for Protected Audience API ran from Chrome 101 to 115. ○ Attribution Reporting API: The origin trial for Attribution Reporting API ran from Chrome 101 to 115. ○ Shared Storage API: The origin trial for Shared Storage API ran from Chrome 101 to 115. ○ Fenced Frames API: The origin trial for Fenced Frames API ran from Chrome 101 to 115.
September Timeline Updates	• "User Agent Reduction" and "Network State Partitioning" were moved from "In Development" phase to "Launched" phase.

Market Testing Grants

In an effort to encourage market participants to test the Privacy Sandbox APIs, Google [announced on July 18, 2023](#) that it has made grant funding available for engineering and testing-related work to eligible SSP and DSP companies to meaningfully contribute metrics that are material to the CMA review of Privacy Sandbox. Google discussed the plans for funding with the CMA prior to announcing the initiative, and the CMA has reviewed the terms of Google's agreements with grantees. Grantees will undertake their testing in line with the [CMA's guidance to third parties on testing](#), and will submit their results directly to the CMA. Google has been providing regular updates to the CMA on the initiative. The CMA can expect grantees to notify them of their testing plans by the end of 2023, and grantees are required to test for at least 8 consecutive weeks between January 1, 2024, and May 31, 2024. Google will continue to engage with the CMA on the progress of this initiative as it develops.

Taking into account observations made by third parties

As part of its commitments to the CMA, Google has agreed to publicly provide quarterly reports on the stakeholder engagement process for its Privacy Sandbox proposals (see paragraphs 12 and 17(c)(ii) of [the Commitments](#)). These Privacy Sandbox feedback summary reports are generated by aggregating feedback received by Chrome from the various sources as listed in the [feedback overview](#), including but not limited to: GitHub Issues, the feedback form made available on privacysandbox.com, meetings with industry stakeholders, and web standards forums. Chrome welcomes the feedback received from the ecosystem and is actively exploring ways to integrate learnings into design decisions.

Feedback themes are ranked by prevalence per API. This is done by taking an aggregation of the amount of feedback that the Chrome team has received around a given theme and organizing in descending order of quantity. The common feedback themes were identified by reviewing topics of discussion from public meetings (W3C, PatCG, IETF), direct feedback, GitHub, and commonly asked questions surfacing through Google's internal teams and public forms.

More specifically, meeting minutes for web standards bodies meetings were reviewed and, for direct feedback, Google's records of 1:1 stakeholder meetings, emails received by individual engineers, the API mailing list, and the public feedback form were considered. Google then coordinated between the teams involved in these various outreach activities to determine the relative prevalence of the themes emerging in relation to each API.

The explanations of Chrome's responses to feedback were developed from published FAQs, actual responses made to issues raised by stakeholders, and determining a position specifically for the purposes of this public reporting exercise. Reflecting the current focus

of development and testing, questions and feedback were received in particular with respect to Topics, Fledge and Attribution Reporting APIs and technologies.

Feedback received recently may not yet have a considered Chrome response.

Glossary of acronyms.

CHIPS - [Cookies Having Independent Partitioned State](#)

DSP - Demand-side Platform

FedCM - [Federated Credential Management](#)

IAB - [Interactive Advertising Bureau](#)

IDP - Identity Provider

IETF - [Internet Engineering Task Force](#)

IP - Internet Protocol address

openRTB - [Real-time bidding](#)

OT - [Origin Trial](#)

PatCG - [Private Advertising Technology Community Group](#)

RP - Relying Party

RWS - [Related Website Sets](#) (formerly First-Party Sets)

SSP - Supply-side Platform

UA - [User-Agent string](#)

UA-CH - [User-Agent Client Hints](#)

W3C - [World Wide Web Consortium](#)

WIPB - [Willful IP Blindness](#)

General feedback, no specific API/Technology

Feedback Theme	Summary	Chrome Response
Ecosystem readiness	SSPs highlighted a concern with publishers not being ready and not doing the required deployment work	Privacy Sandbox has outreach focused specifically on educating publishers, which includes dedicated webinars and meetings with both publishers and SSPs present to drive deployment work.
3PCD	Concerns over 3PCD ramp up in Q4 2023 due to industry tech blackout	The timeline for the Privacy Sandbox has been discussed with the CMA, with sequencing leading to a second-half of 2024 readiness. Privacy Sandbox will publish more detailed information on the sequencing of ramping up 3PCD. Under the Commitments, 3PCD is subject to the CMA's competition concerns being addressed.
Google Ad Manager (GAM)	Google Ad Manager refuses to expose the API surface making testing difficult	Response provided by Google Ad Manager: For the reasons explained below , GAM's plans for its Protected Audience API integration do not include supporting Google's publisher ad server without control of the top-level auction.
Google Ad Manager	Google Ad Manager has a secret floor price that is only exposed to AdX/Open Bidding SSPs.	Google Ad Manager's public documentation says that the winner of the contextual auction is passed to the top level scoring logic and not to any component auction, including AdX or Open Bidding. Furthermore that documentation says of the top level scoring logic: "Ad Manager will compare the winning bid of each component auction, including Ad Manager's own component auction for interest group bids of its buyers, as well as the best contextual ad (which is selected via dynamic allocation), and will serve the ad with the highest bid."
Google Ad Manager	Google Ads products should be subject to the same rules as third-parties' ads products.	Google Ads products are already subjected to the same rules as third parties.
Chrome-facilitated testing	Add labels for browsers not in A or B	We are not considering doing so at this time, as our investigation has found that adding non-experiment labels may complicate privacy concerns around traffic in incognito mode.
Advertising agency	Can agencies or companies without JavaScript on websites use Privacy Sandbox APIs?	Anyone can call the Privacy Sandbox APIs. If an agency or anyone else wishes to build technologies directly on the APIs they can. Client-side APIs require integrating with the client, just as cookies do. Many of the APIs, like cookies, also have an HTTP header interface. We've already seen one ad industry

		framework, Prebid, build client-side integrations with the APIs. Other organizations could do the same.
Client-side Solutions	Why is Google adopting client-side solutions for Privacy Sandbox when an engineer has previously expressed concern on the scalability of such solutions in 2012?	Privacy-enhancing technology (PET) as a field of study has evolved significantly since 2012 and, with it, commercially viable applications. At the core of Privacy Sandbox are combinations of PETs which wouldn't have been feasible over a decade ago. In addition, personal computing power has increased, as have consumer expectations of browsers and regulatory expectations of privacy.
Machine Learning	What is Google's planned usage of Privacy Sandbox for machine learning purposes?	Much of the ad tech ecosystem uses machine learning today and we do not expect that to change. Privacy Sandbox does not prevent ad tech companies or anyone else from continuing to use machine learning. Nor does Privacy Sandbox require that companies integrating with its APIs use machine learning. It is reasonable to expect that companies will continue building products and services in ways that meet the needs of their customers, whether that includes machine learning or not. Any machine learning that Privacy Sandbox integrators do build will obviously be known to them and thus not be obscured to them.
Data verification	How can companies verify that the data they receive from using the Privacy Sandbox is accurate and is Google willing to be reviewed via an entity such as the Media Ratings Council (MRC)?	Privacy Sandbox APIs are built within the open-source platform that powers Chrome. The portions of the APIs meant to run in Trusted Execution Environments are also open source and auditable. Anyone who wants to inspect the code can, including MRC.
(Also reported in previous quarters) Production Support	What is the process in place for Chrome to support Privacy Sandbox technical issues and escalations affecting the ecosystem?	Google provides a range of channels to allow ad techs to report technical issues and enable any necessary escalations to resolve such issues. In addition, Chrome expects to further build and scale a process to resolve technical issues and escalations affecting the health of the ecosystem. Chrome is committed to ensuring resources for this effort. Please see our developer post for more information on the public and private forums for feedback and escalation.
Chrome-facilitated testing modes	More information about the timelines and exact implementations for the Chrome-facilitated testing modes.	We have shared a blogpost about testing modes and are working to share more information soon. We are welcoming suggestions for what size the testing mode labels should be here.

Integration with other industry Standards	Will the Privacy Sandbox APIs connect to either/both TCF V2.* and Consent Mode?	We do not have plans to integrate Privacy Sandbox APIs directly with TCF v2 or Consent Mode. However, companies and industry trade groups are welcome to adapt their products and frameworks to work in conjunction with Privacy Sandbox APIs. For example, with frameworks like TCF each participant must determine its own compliance approach based on the TCF signal it receives and the associated TCF policies. We expect companies to determine when and how to use various functionality our Privacy Sandbox building blocks offer.
---	---	---

Enrollment & Attestation

Feedback Theme	Summary	Chrome Response
Restriction	Enrollment process means Google can decide which company in the ecosystem is allowed to use Privacy Sandbox APIs.	The Enrollment and Attestation process essentially entails verification of the entity (eg. the entity has a DUNs number, can provide a link to a privacy policy etc) and makes the public attestation a requirement for calling the APIs. Entities that can successfully fulfill the enrollment requirements will be validated. For companies that do not have a DUNs, we are providing an expedited, complimentary process with Dun & Bradstreet to acquire one. The objective is to enhance privacy protections of the APIs (by the measures just mentioned) and also to add a layer of transparency to the PS APIs, so interested parties can better understand who is using which API and what attestations they are making. We are open to further industry feedback on this issue, which has already been used to shape the process.
Re-enrollment Overhead	Attestation file expires every 12 months and requires websites to re-enroll.	We've heard feedback from the ecosystem and amended our approach accordingly. This means that files will no longer expire after 12 months or any set period of time. We are updating our enrollment developer guide with additional context.
Attestation file	How is the attestation file used?	All companies calling relevance and measurement APIs will be required by the enforcement deadline to upload the attestation file on their site and keep it for public view as long as you are intending to continue calling the APIs.

		Websites could expect approximately one request per hour from Privacy Sandbox, and other potentially entities may query as well. This will be conducted via the enrollment system's own mechanism to query enrolled entities' servers and ensure the attestation file is valid. Attestations will be included in Transparency Reports and viewable by the general public. We expect companies to act in accordance with their stated attestations, as will the rest of the ecosystem and relevant regulatory bodies.
Enrollment	Is enrollment per site or per origin?	Enrollment is at the site-level.

Show Relevant Content & Ads

Topics

Feedback Theme	Summary	Chrome Response
Performance	Performance concerns on the impact of Topics opt-in rate in the European Economic Area	We would suggest to concerned stakeholders to contact your relevant Data Protection Authority about this issue. They are best-placed to address such concerns and influence whether applications of privacy-enhancing technologies are incentivized by laws or instead treated like tracking, requiring the same approaches to consent. The latter could result in APIs like those in Privacy Sandbox not being available as often.
Enrollment	Do downstream bidders need to enroll in Topics API to use Topics signals from upstream SSPs?	The downstream receivers of topics beyond the initial Topics API caller do not need to be enrolled, though many are likely to be enrolled for other API usage. A list of Privacy Sandbox enrollees will be provided programmatically as part of the program's transparency efforts, which would allow an interested caller of the Topics API to check if the recipient they are sending a topic to is enrolled, if the caller should want to.
Topics filtering	Request to apply another caller's filtering to the topics that they retrieve on the page, in order to only share what buyers are eligible to retrieve.	We are considering this request and welcome additional feedback from the ecosystem.

Site exclusion	Exclude websites from contributing to a user's Topics.	Topics are not called by default. It's important to note that no page content is taken into account when topics are selected, and all topics are curated to make sure they are not sensitive. A website can also restrict their site from being included in topic calculation via the following permission policy header: Permissions-Policy: browsing-topics=()
Topics observation	Allow publishers to give permissions for Chrome to classify topics based on page content (e.g., head, body).	We previously considered offering functionality to classify sites into topics based on page content, and made the decision not to move forward based on privacy and security concerns. This proposal may mitigate some of those concerns, but it's unclear to what extent. Due to the upcoming CMA experiment period, we don't expect this change to occur before 3PCD. We welcome additional feedback here .
Topics observation	Provide more fine-grained permission policies for publishers.	As explained in further detail below , providing more fine-grained permission policies for publishers would enable publisher sites to negatively impact the utility of the Topics API for the ecosystem as a whole, without it negatively impacting the utility of the Topics API for the site itself. Please see this GitHub issue for a more detailed discussion of the topic.
Medical/Health Topics	Why does the Topics taxonomy not cover topics in Medical/Health categories?	Medical and health categories are considered sensitive topics and thus excluded from the Topics taxonomy.
Topics retrieval	Faster way for DSPs to get Topics without fetching using headers.	The header methods are more performant and less costly than creating a cross-origin iframe and making a document.browsingTopics() call from it. (A cross-origin iframe must be used for the call, because the top-level context to observe a topic must match the context from which topics are accessed.) This was discussed in detail here .
Topics retrieval	Requests to support passing Topics via headers on cross-origin script tag requests.	From a security perspective, this isn't possible. Each document and its execution environment are associated with a single origin— that of the document. Third-party subresources loaded and executed within that same environment are considered to be owned by the origin of the document. This is to prevent unconsented data leakage from one origin to another. An alternative is to provide a browsingTopics attribute on <script> tags. This should be clean from a security perspective, and not add additional latency. We are open to feedback from interested parties.
Awareness	Improve public awareness of	We've engaged with the stakeholder who provided this

	Topics API and how the API will be used.	feedback and this issue was resolved on GitHub here. Going forward, we'll continue supporting ecosystem understanding of the API and we look forward to hearing views from stakeholders. In the meantime, we suggest stakeholders wishing to know more about the Topics API to familiarize themselves with the documentation in the Chrome developer guide here.
Notification	Notification to alert user when their Topics are being observed by a website.	We addressed this feedback in Github here . Users can learn more about Topics controls in the Chrome help center here .
Machine Learning	How ML can be used to infer user Topics?	We are discussing this issue and welcome additional feedback here .
Usefulness for different types of stakeholders	Smaller ad tech companies may not be able to observe Topics due to the way browsers calculate them.	Only ad techs that observed the user visit a page about the topic in question within the past three weeks will receive a topic. If the ad tech did not call the API in the previous three weeks for that user on a site about that topic, then the returned value will be empty. This feature means that ad techs whose services are used by a larger number of site owners, and therefore have more opportunities to observe a site visit by a given user, may receive more topics than other ad techs. This feature is essential for the privacy protections of the API as it limits the availability of information about a user to only those parties who are already able to observe the same underlying information (currently via 3PCs).
XHR Request	When will Topics inclusion in XMLHttpRequest (XHR) requests be deprecated?	As Chrome announced in August 2023, Chrome began deprecating support for XHR when transitioning from Origin Trial to General Availability. As the ramp up of Topics progressed, XHR support was only included for users for whom the OT features were enabled and was fully deprecated when the individual OT experiment groups were merged. If you were using Topics with XHR, your sites will not break. The topics just won't be added to your XHR request headers. We recommend that you either transition to Fetch for your request, use the iframe attribute, or the JavaScript API to retrieve topics. Fetch is supported by all modern browsers, but not Internet Explorer or Opera Mini.
Taxonomy and classifier update	More information on the Topics taxonomy and	Our response remains unchanged from Q2:

process	classifier release cadence and how companies can prepare for such updates.	As shared in the recent blog post, we expect the taxonomy to evolve over time, and for governance of the taxonomy to eventually transition to an external party representing stakeholders from across the industry. We also shared the ramp-up plan in the topics-announce group.
Abuse	Potential attack via redirect chain.	We are considering this issue here and welcome additional feedback.
Publisher Inventory Types	What types of publisher inventory will Protected Audience and Topics testing support?	Neither Protected Audience nor Topics are inherently restrictive in terms of the types of inventory they can be used on.
Ramp up time	Recommend no ramp-up time for new taxonomies to get to 100%.	Following this feedback request from the ecosystem and discussion during PATCG meetings, we have announced our plan for the rollout of the new taxonomy here .

Protected Audience API (formerly FLEDGE)

Feedback Theme	Summary	Chrome Response
Top-Level Auctions	Ability to use Google's publisher ad server without also giving Google Ad Manager control of the top-level Protected Audience API auction.	Response provided by Google Ad Manager: GAM's plans for the PA API do not include supporting Google's publisher ad server without the control of the top-level PA auction, for the following reasons. In order to properly serve our customers in the publisher ad serving market, Google's publisher ad server needs to retain control of the top-level PA auction. As a publisher ad server, our role is to provide publishers forecasting so they can negotiate direct sold campaigns without overbooking, and to pace and deliver their direct reservations optimally. Doing this requires running the final auction to compare all eligible direct and indirect demand. Forecasting and pacing are core functionality that publishers expect from an ad server. Without accurate forecasting, publishers may end up overselling their inventory which puts their business reputation at risk. Pacing is also critical - as being unable to fulfill reservation contracts with

		advertisers also risks damage to the publisher <> advertiser direct relationship, which could result in significant impact to a publishers business as. In short, therefore, we do not view a publisher ad server's activity of running the top-level PA auction as distinct from the other activities of the publisher ad server.
directFromSellerSignals	"directFromSellerSignals" allows Google Ad Manager to prevent the publisher from seeing the price of its contextual auction.	Chrome response: Information passed into runAdAuction() is not known to come from the seller unless the seller calls runAdAuction() from its own iframe. In a multi-seller auction it becomes impossible to have all sellers create the frame calling runAdAuction(). directFromSellerSignals addressed this issue by loading content from a subresource bundle loaded from a seller's origin. This ensures the authenticity and integrity of information passed into an auction from the seller—auctions configurations cannot be manipulated. If publishers want to use Protected Audience API to understand any of the information their technology providers are passing into Protected Audience auctions they can ask those technology providers for this functionality. Response provided by Google Ad Manager: We have maintained a strong focus on auction fairness for years, including our promise that no price from any of a publisher's non-guaranteed advertising sources, including non-guaranteed line item prices, will be shared with another buyer before they bid in the auction, which we then later reaffirmed in our commitments to the French Competition Authority. For PA auction, we intend to keep our promise by leveraging directFromSellerSignals, and not share the bid of any auction participant with any other auction participant prior to completion of the auction in multi-seller auctions. To be clear, we won't share the price of the contextual auction with our own component auction either, as explained in this update.
Information Exposure	Sensitive business logics and contractual details may be exposed by the browser.	The person using a web browser can see everything that is happening in the browser. When an ad auction happens inside the browser, it is true that the person whose browser it is could watch

		that auction take place, including seeing how much different parties choose to bid. Since a browser is the user's agent, we do not think it is possible or desirable to try to change this. Only the person using the browser has visibility into these operations, however; an on-device auction run using the Protected Audience API is not observable to any servers, including Google's.
PerBuyerExperimentGroupId	Current value range of PerBuyerExperimentGroupId could allow buyers to correlate the contextual data with the trusted server request.	Using the Protected Audience API in this way is inconsistent with Privacy Sandbox's mandatory attestation that API users will not try to circumvent the Privacy Sandbox's protections. In the future, the requirement that key-value servers run in trusted execution environments (TEEs) will provide technical protection against this attack.
Same-origin policy	Relax the same-origin policy to allow for subdomains.	We are considering this request and welcome additional feedback from the ecosystem here .
API versioning	Request for versioning and release notes for changes to the Protected Audience API.	We are considering this request and welcome additional feedback from the ecosystem here .
Multi-SSP Auctions	Allow top-level auction signals to perform JSON merges with component signal auctionSignals.	We are considering this request and welcome additional feedback from the ecosystem here .
Bid limit	Increase the limit on the number of ad components entering the bid from 20 to 40.	We are considering this request and welcome additional feedback from the ecosystem on why this would be useful here .
(Also reported in previous quarters) Performance of Protected Audience Auctions	Report from testers that Protected Audience auctions have high latency.	On questions of latency, the Protected Audience API has generally followed the existing standard paradigm of building controls that let sellers decide how much time and resources the bidders can consume, and building tools that let buyers decide how to best use the resources available to them. These controls and tools are generally available today, but their full benefit will only be realized after adoption by buyers and sellers. In addition, Chrome continues to work on a variety of infrastructure improvements to auction speed (e.g. crrev.com/1190815, crrev.com/1199839, crrev.com/1201837, crrev.com/1198339, crrev.com/1197323). We invite feedback on both halves of this latency effort: new tools that buyers and sellers would find useful, and reports of observed bottlenecks that Chrome engineers should investigate.

Buy-side filtering	Add support for buy-side filtering based on interest groups.	We have suggested several ways in which SSPs and DSPs could change their designs to handle this: • Moving some work into the DSP's Key/Value server. • SSPs creating some contextual signals and giving those to DSPs. • SSPs caching contextual signals for DSPs.
Publisher Interest Group Control	Support for publishers seeking to 'delegate' the use of publisher-created interest groups.	We have engaged in discussions with many parties about the request. We believe that all such use cases involved in 'delegating' the publisher-created interest groups can be accommodated now, and furthermore that we should build additional support to make some use cases flow more smoothly in the future.
(Also reported in Q2) Trusted Execution Environments	Support for Trusted Execution Environments (TEE) in non-public cloud environments.	Our response is similar to previous quarters: While we are continuing to explore support for options beyond public cloud-based solutions, we have no current plans to support on-premise TEEs. At this stage, given Privacy Sandbox security requirements and the significant challenges presented by on-premise deployments, we believe that continuing to expand and improve cloud-based deployments (for example, supporting Google Cloud in addition to AWS) is the most beneficial for the ecosystem. However, we welcome additional feedback on why such a requirement is necessary and feasible given the privacy and security constraints.
Trusted Execution Environment	Components in the TEE serving path, such as the load balancer, can observe all the traffic and have information of the IP address of each request.	Currently IP address is passed as a metadata in request headers to untrusted seller's ad service in the case of both Bidding and Auction and on-device Protected Audience auctions. See here for more information. In the long term, we plan to proxy ad tech and tracker traffic through an IP Proxy, which will prevent components from observing all the traffic in the serving path.
Time-to-Live (TTL)	Will the time-to-live (TTL) before services have to request new keys be set or is it intended to be flexible (or dynamic)?	The TTL is generally static. Currently, the TTL for the public is 8 days, and the rotation happens every 7 days; the TTL is also the same for private keys in the case of the Aggregation Service. In case of Bidding and Auction services, private and public keys are fetched every N hours in the non-request path and cached in-memory, so that there is no more than an N-hour delay between

		keys rotating and servers picking up these keys. The 1-day buffer between key rotation and expiry is to ensure that even if the key generation fails, the services can continue operating. We are considering extending the TTL to be more resilient for outages. In case of a key leak, we plan to manually force key generation and invalidate keys sooner. Note that public keys are cached on the clients, currently for 24 hours, again to ensure that in case of coordinator outage, the services can still operate.
Traffic Shaping	Traffic Shaping support for Bidding and Auction Services.	Buyers can indicate, based on Publisher first party data or contextual data, demand for Protected Audience auctions. Sellers can do similar determinations as well in the seller's ad server / Ad Exchange server. The models can be trained on 1P data and any aggregate reports from Protected Audience auctions. Sellers can use this information to avoid sending requests to Bidding and Auction servers when there is no demand for Protected Audience auctions. We believe this can be an effective way to shape traffic.
Component Auction	What top level auctionSignals are shared with Component sellers?	Buyers in a component auction only receive signals from the component seller. We are looking to share documentation around the overall sequence of a combined auction with header bidding and Protected Audience auction soon.
Video Rendering	Support for video rendering using Protected Audience and Fenced Frames.	Protected Audience API supports video rendering using a mechanism that relies on iframes. However, we haven't yet designed a solution that is compatible with Fenced Frames, and this is one of the reasons we had decided to push back Fenced Frames enforcement to 2026. That means if a partner does decide to enforce Fenced Frames now, the support for video would be lacking for that partner.
Frequency capping	(Also reported in previous quarters) Per-user frequency controls within a campaign and ad group.	Our response is unchanged from the previous reports: Protected Audience will support frequency capping for on-device auctions and contextual / branding campaigns as well. Shared storage and site-specific caps can also be used for additional frequency capping controls.
Ad Preferences	Does Protected Audience provide a way to opt-out / blocklist by	There are several ways for users to block access to the Protected Audience API and other Privacy

	advertiser sites or a way to leave all interest groups from the same owner?	Sandbox features, as listed here .
Same-origin policy for source url of bidding and auction scripts	Relax the requirement that all fields that specify URLs for loading scripts or JSON must be same-origin with the owner.	We are currently considering this request and welcome additional feedback from the ecosystem here .
forDebuggingOnly	Potential for forDebuggingOnly.reportAdAuctionWin to be misused if it remains post 3PCD.	Over the past years we have been receiving feedback from the ecosystem regarding functionality gaps in Protected Audience once third-party cookies are deprecated, and we are working to formulate a plan to support them post 3PCD without compromising on the goals of Privacy Sandbox. We welcome any additional suggestions and feedback on missing functionality that the ecosystem would like to see.
Multiple Interest Groups	Use multiple interest groups in the same bid.	This is not supported in Protected Audience API today, as it would result in a change to the underlying privacy model. We welcome additional discussion here .
On-device auctions	Will Chrome on Android support on-device Protected Audience auctions?	Yes, on-device auctions will be supported in Chrome on Android, as indicated here .
(Reported in Q2 2023) Click related data	Add click-related data to browserSignals.	We continue to evaluate this feature request and welcome additional feedback on why this should be prioritized here .
Trusted Execution Environment providers	Are there material differences in the Trusted Execution Environment offerings of different cloud providers?	We are not aware of any major differences, but we recommend the ecosystem review the public deployment guides to see which solution best suits their needs. Google Cloud – here. AWS – here.
(Reported in previous quarters) Support for negative Interest Group targeting	An API to support negative interest group targeting: showing ads only if a user does not belong to an interest group.	We are looking into implementing this feature and are discussing the request here .
Content Violation	Support features that allow users to report bad ads served by Protected Audience API in Fenced	We believe that the existing Fenced Frame Ads Reporting mechanism offers good options for

	Frames.	ad techs who want a user-generated "Bad Ads" reporting flow. This would allow bad ads reporting in a way essentially unchanged from the industry standard today. We welcome additional feature requests if any gaps remain, including during the time after third-party cookie removal but before Fenced Frame rendering becomes widespread.
Private Aggregation API Reporting	How to calculate time the user has spent in that interest group?	In Chrome M116+ you should be able to use recency as defined here .
K-Anonymity server	More information on K-Anonymity server.	We shared more information on K-Anonymity servers here and welcome additional feedback.
Dynamic Creative URLs	Support for creative URLs without pre-declaration while still respecting K-anonymity.	We are discussing this feature request and welcome additional feedback on why this should be prioritized here .
k-anonymity requirement	Will k-anonymity requirement on Interest Group updates be re-introduced?	We don't anticipate changes to the position stated in this GitHub post. As announced in that post, we decided to remove the k-anonymity requirement on Protected Audience interest group updates, which does not have a significant impact on the API's overall privacy protections, and we plan to consider other potential more direct protections (e.g. IP address privacy or a trusted update server) at a later date when the related technologies are more developed, deployed and adopted.
Bidding & Auction Services Beta Testing	When will Bidding & Auction Services Beta testing begin?	As stated here , the first phase of Bidding and Auction Services testing begins in November 2023.
Roadblocking	Request to support Creative coordination for Ad Networks (SSP and DSP are in the same company or properties).	We appreciate the feedback for this use case and we're looking to understand whether more ad techs are interested in seeing this supported. We welcome additional feedback here .
Native Advertising	Fenced Frame support for Native Advertising.	We are considering supporting the use case and are discussing possible workarounds and solutions here .
K-anonymity	How can I maximize interest group ads that meet k -anon thresholds?	We have shared some tactical guidance on this topic here .
POST support	Support for sending auction data via POST requests.	We are evaluating this feature request and welcome additional GitHub issue submissions here on why this should be prioritized.
Reporting granularity	What is the reporting granularity of Fenced Frame Ad Reporting	Current design does not allow capturing product ID and/or position as this may compromise user

	with Ads Composed of Multiple Pieces?	privacy. Only the reserved.top_navigation can be invoked, which would be sent when there is a user activation (e.g. click) on the ad component fenced frame, which results in a top-level navigation.
Ad Auction	Can SSP participating in a component auction trigger another component auction itself?	A componentSeller cannot also include componentAuctions. The multi-seller auction only has two levels: 1. The component auctions in parallel. 2. The top-level auction (where the winning ad from each componentAuction competes).
Bidding & Auction Services availability	Will Bidding & Auction be available during the Chrome facilitated testing phase?	Bidding and Auction Server will not be available during the Chrome facilitated testing phase.
Bidding signals	Allow browsers to request and delete bidding signals.	We are discussing this request and welcome additional feedback on why this should be prioritized here .
generateBid()	Ability to update interestGroup's userBiddingSignals through updateURL.	We are considering this proposal and welcome additional feedback and discussion here .
Publisher Inventory Types	What types of publisher inventory will be supported by Protected Audience and TOPICS testing?	Neither Protected Audience nor Topics are inherently restrictive in terms of the types of inventory they can be used on.
Server to Server integration	Is direct integration between SSP and DSP required for Protected Audience?	Direct integration between SSP and DSP is not required if DSP does not need to process contextual signals in its own server in order to pass that processed information into its on-device bidding function.
bid_currency field in B&A	Support for bid_currency field in Bidding and Auction Service.	B&A doesn't support bid_currency yet, although we plan to support that by the end of January 2024. See timeline here .
perBuyerSignals	Is there a size limit for perBuyerSignals?	There is no limit on the number of per-buyer signals, but sending too much data may have detrimental effects on the browser's performance.
Cross-site use cases	Can we use Protected Audience API interest groups across multiple websites?	Protected Audience is not designed for such use cases, as explained here .
Interest Group HTTP Requests	Include Interest Group Blob in the HTTP headers.	We are considering this request and welcome more feedback on this request here .
Ad Quality control	Loss of ad quality control related on cross-site information.	We are considering this feedback and welcome additional feedback here .
Chrome DevTools	Outgoing Protected Audience network requests should be visible	We are working on enabling this functionality in the network tab and welcome additional feedback

	in the Chrome Developer Tools Network Tab.	on why this should be prioritized here .
Trusted Execution Environment	When will the details on which metrics are privacy impacting (and their degree) be added to the explainer on Trusted Execution Environment monitoring?	We are in the process of updating the explainer with this information. The updated explainer will be available by November 2023.
directFromSellerSignals	Why is directFromSellerSignals not packaged as a web bundle?	We shared the rationale for this decision here .
Impression delegation	Is there any viable way to do impression delegation where the outcome of an interest group being selected is yet another targeting action?	Multiple nested auctions are not compatible with our privacy goals for two reasons. First, when the winner of an auction renders inside a Fenced Frame, our privacy goals for Protected Audience include the resulting creative rendering without knowledge of the context: surrounding page's URL or first-party cookie are a privacy violation. In that environment, a nested auction is not viable. Second, the Protected Audience model says that each auction's winner should be based on data from just one additional site. Nested auctions would be a way to compound that, resulting in the possibility of choosing ads based on a many-site profile.
Data at Rest' criterion	Explain further the 'Data at Rest' criterion in the Key/Value service trust model.	Data in the Key Value Service is loaded into memory and served from there rather than doing any read-through caching.
Buyer Data Signal	Is there a defined size limit for the buyer_data signals received from the DSPs?	There are currently no browser imposed limits for buyer_data signals received from DSPs.

Measuring Digital Ads

Attribution Reporting (and other APIs)

Feedback Theme	Summary	Chrome Response
Cross device	Plan for cross-device support for Attribution Reporting API.	Cross-device presents new privacy challenges on top of 3PC and also adds technology distribution challenges given the range of devices and platforms a user might use. We are exploring potential solutions, but we are focused on the critical use

		cases currently supported by ARA and do not have plans to introduce cross-device support before the removal of third-party cookies.
(Also reported in previous quarters) Trigger Data Size	Why is the trigger data size limited to 3 bits?	The size is limited to 3 bits and 8 distinct values to ensure that the amount of cross-site/context information about a user is limited. We welcome ecosystem players to submit feedback on whether the current parametrization for event level reporting is sufficient here .
Conversion funnel	Report multiple domains that were used in conversion.	This use case is possible since the addition of multiple destinations – see here . We welcome additional feedback here .
Same Domain in different country support	Does Attribution Reporting work with websites that have the same Domain but multiple country TLDs?	This issue has been discussed and resolved with the stakeholder that raised the question. If an ad tech needs to use multiple country TLDs they will need to have multiple enrollments, with one for each country TLD.
Protected Audience and Attribution Reporting	Can ad techs access both view-through conversions for Protected Audience auctions as well as click-through conversions for Attribution Reporting?	Yes, Privacy Sandbox should support both VTCs and CTCs within Protected Audience.
Aggregatable report delays	Reduce aggregatable report delays further.	We have heard recent feedback regarding this and have shared ideas here . We welcome additional feedback from the ecosystem.
Aggregatable report delays	Reducing delays via introducing server mediation.	We are considering this proposal and welcome additional feedback here .
Event level report delays	Reduce event-level report delays.	The full flexible-event level proposal – see here – can reduce event-level reporting delays down to 1 hour with a noise tradeoff.
Source reporting origin per source	Limitation of max source reporting origins per source reporting site prevents ad techs from registering sources from different reporting origin for a single publisher origin.	This has been discussed with the stakeholder that raised the issue and a potential solution of using 1 reporting origin per source-reporting site is being tested before trying other potential solutions involving redirects. We are open to any additional ecosystem feedback regarding this limit as well.
Issue reporting	How to report errors or issues with Attribution Reporting API to Chrome.	Currently we recommend ad techs report any Attribution Reporting API errors they may be facing as an Issue on Github. If they are facing a Chrome related issue we recommend them creating a Chromium bug. Links for how/where to flag any

		issues can be found here .
Deduplication	How to deduplicate conversions across different pipelines and devices.	Deduplicating across devices and measurement pipelines is a known and current challenge that ad techs also face today with 3PCs. With the Attribution Reporting API, ad techs can decide when to register specific conversions and add specific metadata to indicate which measurement pipelines they have used to track the conversions (i.e. part of the aggregation key), which can be compared against other measurement pipelines. We are open to any additional ecosystem feedback regarding this.
Deduplication and Priority	Request to have priority first before deduplication.	We are considering this request and welcome additional feedback here .
Anti-fraud	Risk of malicious user tampering the event-level data.	Report verification does not work for event-level reporting for the reasons described here .
Conversion type	How to differentiate between view through and navigation in Attribution Reporting.	We have the following built in filtering option: "source_type" (additional details can be found here).

Aggregation Service

Feedback Theme	Summary	Chrome Response
Budget recovery	Some adtechs have requested the ability to reprocess reports in cases where there are failures/errors/deletions of their reports.	The team is exploring ways to address this in a privacy-preserving way.
Site enrollment	Multiple ad techs have requested support for processing multiple origins in the same account for use cases such as splitting data by Geo, advertiser. This behavior is also expected by ad techs given that the client API enrollment is now site-based (and not origin based). Migration from origin to site enrollment streamlines the ad tech onboarding process via consistency with the client enrollment process.	We will be launching migration from origin enrollment to site enrollment for the Aggregation Service soon and welcome feedback from the ecosystem here .
Release &	Release and depreciation schedule	We have recently published a proposal for the

Deprecation Plan	for Aggregation Service features and patches published. The goal of the plan is to give adtechs visibility into our release policies to enable them to prepare for upcoming releases and deprecations, and ensure they run stable and secure versions of services.	Aggregation Service release and deprecation plan here and welcome additional feedback here .
Coordinators	What happens if the coordinators go down on aggregation service?	Both coordinators need to be fully available for the system to function correctly. Short unavailability is accommodated with retries in our client libraries; longer unavailability of either of the two coordinators will have aggregation jobs fail. Jobs can be rerun if the budget for privacy isn't consumed yet. In the case where any service failure led to budget consumption without a summary report written to ad tech storage, we currently recommend they use debug reports to retrieve results using the local testing tool here. We are also working on features to allow for budget recovery in the case of failures so adtechs can rerun their jobs.

Private Aggregation API

Feedback Theme	Summary	Chrome Response
Blob Url	Request to support Blob Url in Shared Storage.	Support for Blob Url has been added in Chrome M116.

Limit Covert Tracking

User Agent Reduction/User Agent Client Hints

Feedback Theme	Summary	Chrome Response

JavaScript API	Availability of the User Agent Client Hints JavaScript API.	There are no plans to remove this functionality as it is our core solution for partners who want to actively access the high-entropy data beyond what is available by default in the frozen and reduced UA.
Device/Form Factor information	Ability for websites to understand input, output, and other information the device visiting the website can support.	We have added support for this request here following feedback from the ecosystem.

IP Protection (formerly Gnatcatcher)

Feedback Theme	Summary	Chrome Response
Eligible Third Party Traffic	What is "eligible third-party traffic" referring to in the explainer?	We understand the importance of this question and are actively working to identify which third-party traffic will be eligible and which will not. We welcome feedback on this topic.
Network Traffic Audits	Support for enterprises to perform network traffic audits for their networks.	Only third-party traffic embedded in first-party sites will be affected, which should limit the amount of traffic that requires filtering. Additionally, we plan to give users the option of whether or not to use IP Protection, and for enterprise-controlled Chrome, there will be enterprise policies to disable IP Protection. Finally, we're exploring what controls (if any) will be provided to network operators to disable IP Protection. We welcome feedback on this topic.
Access control	IP Protection may impact web services that use IP addresses for access control.	We understand the importance of anti-fraud use cases and the possible impact to those use cases. We are seeking ecosystem feedback on how we can better support anti-fraud use cases that typically have relied on IP addresses.
Communication between the 2-Hop proxies	How to ensure there is no information between proxies.	We are in the process of designing the proxy interactions. Our goal is to minimize the chances for such information sharing via business, process, and technical means.
Non-Google Authentications	Support for Non-Google Authentications.	We plan to publish more details about account authentication in the future, though we have shared some initial considerations already here .
Tracker classification	How will IP Protection determine what constitutes a tracker and its variants?	We understand the importance of this question and are actively working to identify which third-party traffic will be eligible and which will

		not. We welcome feedback on this topic.
Analytics	IP Protection may impact the accuracy of analytics services.	We are looking to understand the impact of IP Protection further and welcome additional feedback and examples from the ecosystem.
Proxy	If a user is using proxy or has manually defined a proxy, how will IP Mask work in this case?	We are looking to understand the impact that IP Protection may have on other proxies. We do not have any plans to share at the moment. We welcome feedback on this topic.
Premium offering	Will IP Protection be a paid feature?	IP Protection will be available to Chrome users as part of the core browser experience. It will not be a paid feature.
Proxy server	Will the same proxy servers be used during user sessions?	An HTTP/S connection will use a single pair of proxies and will present a single masked IP address to the origin. Beyond that, there are no hard constraints on different HTTP/S connections having to use the same servers.
Platform support	On which platform will IP Protection be supported?	IP Protection will initially be available on Chrome for Android and Desktop. We continue to evaluate how to expand the protection to other platforms.
Opt-Out	Will users be able to disable IP Protection?	We plan to provide users the choice on whether they want to use IP Protection or not.
Anonymization	What kinds of requests will be anonymized under IP Protection?	HTTP/S and DNS requests to eligible third-party domains are anonymized via the privacy proxies. We will provide additional details in an upcoming explainer on how we will determine which domains will be included. The rest of the traffic (e.g., the rest of the DNS requests or other HTTP/S traffic) is unaffected.
Data Visibility	Network addresses may be accessed during the first hop in IP Protection.	In the two-hop proxy model, the first hop (controlled by Google) only sees the source client IP and a request to connect to the second hop, while the second hop (controlled by an external CDN) only sees a tuple on the first hop (proxy IP + port) and the destination IP. For the response back from the origin, the second hop is able to forward the response to the first hop proxy+port associated with the request and doesn't need to learn anything about the original client IP (and the first hop just returns the response to the client, without learning anything about the destination IP). In this way, the first hop only learns the client IP and the second hop, while the second hop only learns the destination IP.
WebView	Will IP Protection be available to	We do not have any plans to share at the moment,

	Android WebView in the future?	but our vision is to provide this protection as broadly as possible.
--	--------------------------------	--

Bounce Tracking Mitigation

Feedback Theme	Summary	Chrome Response
Interaction Tracking	How are user interactions tracked?	Bounce tracking mitigations track two types of user interactions: 1. User activations as defined by the html spec. These are basically clicks, key presses, touch screen taps, etc. 2. Successful webauthn assertions. These are cases where a user taps a security key or uses a passkey as form of authentication These interactions are associated with the top-level site on pages where they occur. For example, if a user clicks in an embedded iframe the interaction is associated with the top-level site and not the embedded site. The interactions are stored in a database containing the schemeless etld+1 and the time of the interaction. Interactions protect the associated domain from bounce tracking mitigation state deletion for 45 days.
Allowlisted Exemptions	Can domains be exempted?	We are considering this request and we welcome additional feedback from the ecosystem here .

Privacy Budget

No feedback received this quarter.

Strengthen cross-site privacy boundaries

Relative Website Sets (formerly First-Party Sets)

Feedback Theme	Summary	Chrome Response
Centralized Approach	Concern over the centralized repository approach for managing Related Website Sets.	A public, easily accessible repository is key to the design of RWS as it provides accountability for submissions. Third-party cookie functionality is ultimately provided by the use of the Storage Access API or the rSAFor API, with RWS membership providing auto-granted access (as opposed to through prompts with the Storage Access API). We believe that an approach like the RWS submission process is an appropriate requirement for auto-granted third-party cookie access.
Renaming json file	With the change in API name, does the hosted JSON file name need to be changed?	Yes, the submission guidelines have been changed, and the primary domain must serve a JSON file at <code>/.well-known/related-website-set.json</code>. Existing sets in the RWS list do not need to be changed, but if there are modifications submitted to existing sets, the JSON file must be changed.
(Also reported in previous quarters) Domain Limit	Request to expand the number of associated domains	As announced in a blogpost on August 31, we have raised the associated domain limit to five domains following feedback from the ecosystem. We have decided to increase the associated domain limit to five domains (plus one primary domain) which best matches the most comparable implementation offered by another major browser.
Third Party Cookies	Will Related Website Sets only work with third-party cookies disabled?	Related Website Sets will work even when a user has not blocked third-party cookies; but there will be no observable effect since the relevant cookies are available without any need for Related Website Sets and Storage Access API.
Legitimate edits	How does the Related Website Sets repository prevent non-owners from modifying sets?	Per the submission guides , anyone can submit a PR on GitHub to edit the <code>first_party_sets.JSON</code> file. However, if the PR is approved (passes technical validations, etc.), it will be manually

		merged in batches to the canonical FPS list once per week (Tuesdays at 12pm Eastern Time) by Google. If a bad actor tries to modify a set they don't own, it shouldn't be a problem since they won't be able to modify the .well-known files and therefore the validations will fail.
Domain hijacking	Domain hijacking may expose related domain data to unauthorized parties.	This is not possible, as discussed in this GitHub issue.

Fenced Frames API

Feedback Theme	Summary	Chrome Response
Content Violation	Allow users to report suspicious ads.	Suspicious ad reporting is not prevented by Fenced Frames. Users can still interact with the ad and report suspicious ads to the ad tech in the usual way.
Interaction with surrounding sites	Allow interaction with the surrounding/top-level website.	We are looking to understand why this request is necessary and welcome additional feedback from the ecosystem.
Native Advertising	Fenced Frame support for Native Advertising.	We are considering supporting the use case and are discussing possible workarounds and solutions here .

Shared Storage API

Feedback Theme	Summary	Chrome Response
Cross domain	Allow communication across domains for local storage.	This use case is currently not in line with Shared Storage's privacy-preserving output gates but we welcome additional context here as we evolve proposals for non-partitioned storage.
Blob Url	Request to support Blob Url in Shared Storage.	Support for Blob Url has been added in Chrome M116.

CHIPs

No feedback received this quarter.

FedCM

Feedback Theme	Summary	Chrome Response
Third Party Cookies	Is FedCM currently disabled if users enable "Block third-party cookies" in the Chrome settings?	Yes, FedCM is currently disabled. For testing, we recommend that you additionally enable <code>chrome://flags/#fedcm-without-third-party-cookies</code>. We are looking to support FedCM without third-party cookies in the future, as discussed here.

Fight spam and fraud

Private State Token API (and other APIs)

Feedback Theme	Summary	Chrome Response
Token expiration	Once Google Chrome is uninstalled, will the Tokens be lost or will it be cached?	The token will be lost if the user uninstalls Google Chrome.
Token Information	How can issuers keep issued information within the Private State Token private?	Information is always kept private in the token and cannot be unencrypted by external parties that do not have the keys.
Error in demo	Error when trying to run the Private State Token demo.	We have updated the demo and it should be working correctly now.

Google Ads Roadmap for Effectiveness Testing of the Privacy Sandbox Proposals

Google Ads is engaged in integration and testing of the APIs and providing feedback to the CMA and the ecosystem. Google is conscious of the importance of transparency for the ecosystem, so that they can plan their investments and forecast participation in future tests, and as such has included Google Ads' testing plans below:

Protected Audience API for Remarketing:

- In Q4 2023, Google Ads plans to conduct an experiment with the Protected Audience API (individually) for Remarketing on Chrome Desktop and Mobile Web utilizing General Availability traffic from the Google Display Network.

Measurement APIs:

- In Q4 2023, Google Ads envisages publishing guidance on how third-party ad tech could improve Event and Aggregate-API data from the Privacy Sandbox Attribution Reporting API via intelligent configuration.
- In Q4 2023, Google Ads plans to conduct an experiment with the Attribution Reporting API (utilizing both Event-level and Aggregate-level reports) on Chrome Desktop and Mobile Web utilizing General Availability traffic from a subset of Google Owned and Operated properties.
- In Q1 2024, Google Ads plans to continue the experiments with the Attribution Reporting API (utilizing both Event-level and Aggregate-level reports) on Chrome Desktop and Mobile Web utilizing General Availability traffic from an expanded set of Google Owned and Operated properties.

Chrome-facilitated testing:

- In Q1 2024, Google Ads plans to conduct an experiment to test privacy-preserving solutions and Chrome's Privacy Sandbox APIs in combination (Topics, Protected Audience and Attribution Reporting) via [Chrome-facilitated testing](#) on Desktop and Mobile Web with traffic from the Google Display Network. We encourage authorized external parties (Demand Side Platforms aka DSPs and Supply Side Platforms aka SSPs) to participate in this experiment with us.

Google's long term testing timeline, along with registration details for Chrome's Origin Trials and details of the APIs is available at the privacysandbox.com site.

Google's Interactions with the CMA

Efforts to identify and resolve concerns quickly

Paragraph 15 of the Commitments provides for Google to engage with the CMA in an open, constructive and continuous dialogue in relation to the development and implementation of the Privacy Sandbox proposals, in the context of which paragraph 17(a) envisages efforts to identify and resolve concerns quickly.

The intensive discussions between Google and the CMA set out below have focused on ensuring that the CMA is fully informed of developments in the Privacy Sandbox proposals, and of the underlying thinking. Google continues to respond to a continuous sequence of detailed questions in this respect. As part of this, the parties continue to operate a joint process by which the CMA carefully reviews relevant Google announcements before they are published.

CMA concerns

The CMA has not during the relevant period expressed concerns for resolution through the formal routes pursuant to paragraph 17(a)(ii), or notified any such concerns pursuant to paragraph 17(a)(iii) of the Commitments. However, the CMA has continued to raise detailed questions and concerns about how the Privacy Sandbox APIs would address the Development and Implementation Criteria set out in the Commitments, based on its own assessment and reacting stakeholder concerns as set out below, and required Google to respond to these submissions.

Stakeholder concerns

The CMA has shared with Google certain concerns expressed by stakeholders including the following:

3PC Phaseout - The CMA has highlighted to Google a blog post from [RTB House](#) on how to phase out 3PCs and whether this should be a gradual process. The CMA has also shared stakeholder feedback that it would be preferable to define exit criteria rather than having a timeline for 3PCD. Google is still in the process of considering the dynamics of how 3PCD will be facilitated, but in order to minimize the impact on the holiday advertising season, does not intend to phase out 3PCs during the 2024 end-of-year code freeze. To provide reassurance to the ecosystem that this will be a gradual process, Google has updated the hover-over on "Third-Party Cookie Phase Out" on the [timeline](#) on [privacysandbox.com](#) to state "*Chrome will begin gradually phasing out support for third-party cookies in Q3 2024.*" Google has also agreed to make clearer in its public communications that 3PCD is conditional on the CMA's competition concerns being addressed, and that disabling 1% of third-party cookies from Q1 2024 is for the purposes of facilitating testing and is not the

start of 3PCD. 3PCD would only happen after a standstill period of at least 60 days, as set out in the Commitments.

Protected Audience API (formerly known as FLEDGE API)

The CMA has shared that a stakeholder has expressed concern regarding the fact that Protected Audience API auctions using the Trusted Execution Environment (TEE) only supports AWS and Google Cloud Platform (GCP). Services running in the TEE should be deployed on a cloud platform that supports necessary security features. In addition to AWS and GCP, Google expects to support other cloud providers in the future, and Google is open to suggestions for other cloud providers. See [this explainer](#) for more information.

The CMA has shared stakeholder feedback regarding the number of open issues on GitHub relating to the Protected Audience API. As was noted on GitHub, Google has suggested labeling open issues, to allow the team to see which relate to future enhancements, which are editorial, and which may have an impact on the processing model or API shape in ways that can impact future compatibility. We welcome any additional feedback on this topic.

The CMA has shared stakeholder feedback that Privacy Sandbox introduces parallelism of auctions, creating a risk that an advertiser might bid against itself. Google has not received feedback directly about risks from parallelism, and is open to discussing ways for a particular DSP to avoid bidding in a single auction through many different SSPs if this is a desirable feature to add to the API. Interested stakeholders can provide feedback to Chrome directly [here](#), or open an issue on GitHub for public discussion [here](#).

The CMA has shared that a stakeholder has expressed a concern that Google Ad Manager (GAM) has indicated that they intend to throttle randomly and later via machine learning, the Protected Audience API runner. The stakeholder considers that the publisher must be able to override this behavior in either direction, as the model will likely optimize AdX revenue. GAM understands that some market participants might mistakenly have inferred that GAM has plans for random throttling, because GAM currently only enables Protected Audience API auctions for a limited percentage of traffic for testing purposes. However, that's a misunderstanding. GAM does not intend to do any random throttling. As explained in GAM's [help center](#), GAM is in the process of gradually increasing the amount of traffic included in GAM's testing and expect that by the end of 2023, up to 10% of Chrome traffic will be enabled for Protected Audience API testing. GAM also intends to use machine learning models to determine whether to trigger a Protected Audience API auction to optimize for total publisher revenue from all sources (including direct sold reservations, revenue from programmatic auctions from AdX, and revenue from other SSPs a publisher works with).

The CMA has shared stakeholder feedback regarding Google refusing to share information on which DSPs have interest groups associated with a given impression. Google has pointed out that the requested information would constitute a probably-unique user

identifier, so would defeat the tracking prevention goals of Privacy Sandbox. However, as noted [above](#), Google has also suggested several alternate ways in which SSPs and DSPs could change their designs to handle traffic load, including (i) Moving some work into the DSP's Key/Value server; (ii) SSPs creating some contextual signals and giving those to DSPs; and (iii) SSPs caching contextual signals for DSPs.

Topics

The CMA has shared stakeholder feedback that random noise gives a statistical signal/noise advantage to callers that are present on large numbers of sites. This issue was raised by the CMA to Google (see also the corresponding GitHub issue [here](#)) and in December 2022, Chrome introduced a change ensuring that a fixed 5% of returned topics are random, regardless of the caller's presence. However, a caller with less presence, will nonetheless still receive fewer different topics. By design, in order to be a privacy improvement over third-party cookies, the Topics API caller should learn no more than it could have using third-party cookies. This means the API shouldn't inform callers about topics that the caller couldn't have learned for itself using cookies. The topics that a caller could have learned about using cookies, are the topics of the pages that the caller was present on with that same user. This is why the Topics API restricts learning about topics to those callers that have observed the user on pages about those topics.

The CMA has also shared stakeholder feedback that publisher sites should be able to retrieve topics whilst opting out of training, or alternatively third-party callers should also be denied this option. While third-party callers can indeed decide to retrieve topics without observing, third-party callers cannot do this all the time, because the per-caller topic filtering mechanism means that an API caller who never observes topics will also never retrieve topics. In contrast, allowing publisher sites to retrieve topics while opting out of training would enable publisher sites to negatively impact the utility of the Topics API for the ecosystem as a whole, without it negatively impacting the utility of the Topics API for the site itself. Please see [this](#) GitHub issue for a more detailed discussion of the topic.

The CMA has shared stakeholder feedback that there is a persistent data advantage to large, multi-topic sites, and therefore, topics should be page-level rather than site-level. Rather than characterizing this dynamic between large and small, we consider that the relevant distinction is between general-interest, and niche-interest sites. With 3PCs, there is a difference in the value of information contributed by different sites. Niche-interest sites are inconsistent in their value contribution: not all niche-interest sites have commercially-valuable context, and therefore contribute less value. These are the sites which will benefit from the Topics API. We have considered the possibility of page-level rather than site-level classifications, however, there are a number of significant privacy and security concerns with such a classification.

Topics Taxonomy – The CMA has shared that a stakeholder expressed that the Topics taxonomy is twice as granular as Google's current buy side segments, which will have

negative impacts on publishers of all sizes. Google understands that this feedback relates to Google's "affinity" taxonomy which contains 154 categories. However, [Topics was designed to address the interest-based advertising use case](#), which includes affinity and in-market sub-use cases. A more appropriate comparison would therefore include Google's in-market taxonomy, which contains 779 categories. Therefore, with 469 categories, the Topics taxonomy contains approximately *half* as many categories as the in-market and affinity buy-side segments. Nonetheless, one of the key issues that we have considered as we worked on the revised taxonomy is the granularity, which was also a focus of stakeholder feedback in the past two quarters. The CMA shared that some ecosystem participants feel that the taxonomy should not be made any more granular due to privacy concerns, whereas others conversely would encourage more granularity. Google has sought to develop a taxonomy with categories that better match advertiser interests, while maintaining our commitment to exclude potentially sensitive topics and safeguard user privacy.

The CMA has shared stakeholder feedback that the granularity of Topics could undermine a publisher first-party data solution, and that changes to the Topics taxonomy make it more likely for proxies to be created for the niche publishers' audience and be available to purchase from a generic site for free, which will harm the industry. We are conscious of the tradeoffs underlying the sizing and granularity of the Topics taxonomy, and we look forward to continuing to engage with the CMA and ecosystem on this issue. The APIs will help to enable 3PCD which, overall, will increase the value of publisher first-party data. While a more granular Topics taxonomy may indirectly decrease the appeal of other solutions, such as those based on publisher first-party data or those relying on direct deals, as we develop the Topics API our main goal is ensuring that it supports interest based advertising use cases after 3PCD as effectively as possible, for all stakeholders alike. Our belief is that greater utility for Topics will improve competition overall and benefit the ecosystem as a whole.

Latency scale - The CMA has shared feedback that moving processing to the browser could introduce unacceptable latency. Keeping latency to a minimum is a key design goal of the Privacy Sandbox APIs. Our current expectation is that API latency should have minimal impact on a site's Core Web Vitals, as the majority of APIs are called after the initial rendering of the website. We will continue to monitor our other standard browser user experience metrics, like scroll jankiness, for evidence of latency impact due to post-rendering execution (though we have not yet decided the duration of the holdback which enables this type of monitoring). We continue to monitor and make improvements to reduce latency further for each of the APIs, and encourage continued testing and feedback

First-party data - The CMA has shared stakeholder feedback that Google's use of first-party data is not necessarily safer from a data protection perspective than others' use of third-party data, and that Google's Privacy Policy (as updated at the beginning of July 2023, to refer to add Google's Bard service as an example of the AI products powered by AI models trained on publicly accessible data) does not (or did not) specify the use to which

the personal data will be put. Google goes to great lengths to protect the first-party data we hold, as well as putting users in control of their own data. Among other things, Google limits the personal information that's used and saved, provides clear transparency and user controls over data that is collected, uses advanced privacy techniques to keep personal information private, has a strict policy against selling personal data, and has strict privacy protocols that are followed throughout every product's development. For clarification, Google's Privacy Policy has long been transparent about the use of publicly accessible info to train the company's language models that power services like Google Translate; the July update simply added Bard as an example of the services powered by similar models. Further details of how Google protects user data are available [here](#). Interested stakeholders can provide further feedback on data protection directly to Chrome [here](#).

Interoperability and Testing - The CMA has shared stakeholder feedback that Google's APIs will not be interoperable with other browser engines, and that there are limits placed on a publishers' ability to test anywhere other than within Google's systems. As noted in our Q4 2022 progress report, Google's long-term goal remains to create interoperable standards that multiple browsers broadly support and that provide effective, privacy-enhancing solutions for targeting and measurement use cases. Multiple Chrome teams and many third parties are working in various W3C groups such as the [Private Advertising Technology Community Group](#) (PATCG), [Web Platform Incubator Community Group](#) (WICG), [Federated Identity Community Group](#), and others, to identify and work on solutions that are broadly acceptable across many browser engines. While Chrome is open to the adoption of the Privacy Sandbox APIs by other browsers, Chrome is not currently aware of other browsers implementing the Privacy Sandbox APIs as a solution (see [here](#) for further details). For publishers seeking to test the APIs on Chrome, publishers are not required to do so with Google Ads, and Google encourages all participants in the ecosystem to test the APIs.

Status Meetings

The Commitments provide for Google and the CMA to schedule regular meetings at least once a month (before the Removal of Third-Party Cookies), to discuss progress on the Privacy Sandbox proposals. Currently, Google and the CMA typically have one substantial technical meeting a month, updating on progress and addressing an agreed agenda of testing, targeting, measurement, boundaries and user control topics to assist the CMA to carry out the regulatory scrutiny and oversight foreseen in the Commitments, as well as one legal status meeting focusing on legal, procedural, and competition considerations. Google and the CMA collaborate on the agendas for each meeting to ensure that adequate attention is given to each topic. Additional meetings are held to discuss specific issues when the need arises.

In addition to synchronous meetings, Google and the CMA typically engage with each other on at least a weekly basis. These engagements range from emails to formal written responses, and consist of questions and answers, the sharing of information, and the like.

Standstill

Paragraph 21 of the Commitments on notification of concerns during the Standstill is not yet applicable, as Google has not entered the Standstill Period.

Compliance statement

The compliance statement provided for at paragraph 32(a) of the Commitments is attached.



COMPETITION AND MARKETS AUTHORITY
Case 50972 - Privacy Sandbox
Compliance Statement

I, Renée M. DuPree, Director, Competition Compliance of Google LLC confirm that for the three months to 30 September 2023, Google has complied in the preceding three-calendar-month period with the obligations relating to:

- Google's use of data set out in paragraphs 25, 26, and 27 of the Commitments;
- Google's non-discrimination commitments set out in paragraphs 30 and 31 of the Commitments; and
- Google's commitment in relation to anti-circumvention in this respect set out in paragraph 33 of the Commitments.

Any failures to meet the Commitments during this three-calendar-month period were notified to the CMA within five Working Days of Google becoming aware of them and are also listed below for completeness.

Signed.....

Full name.....

Date.....

Breaches (if any) listed on following page for completeness: Not applicable