



Privacy Sandbox Progress Report

Q4 Reporting Period - October to December 2024

Prepared for the CMA, 31 January 2025

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 [developer documentation](#) which provides overviews for the core [private advertising](#) features and [cookie changes](#), along with API implementation and [status information](#). Key updates are shared on [the developer blog](#) along with targeted updates shared to the individual developer mailing lists.

Updated Timing Expectations

In July 2024, Google provided an update on [A New Path for Privacy Sandbox on the Web](#). An overall timeline update is pending as Google remains in ongoing discussions with the CMA and ICO. Google's latest expectations for the timing of the individual Privacy Sandbox proposals are set out in the [Privacy Sandbox Timeline](#).¹ The summary below includes all Q4 2024 updates, covering the period from October 1 to December 31, 2024.

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 Q4 2024 Timeline Updates	
Oct Timeline Updates	• No changes
Nov Timeline Updates	• No changes
Dec Timeline Updates	• No changes

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, PA API and Attribution Reporting APIs and technologies.

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

Glossary of acronyms.

ARA - [Attribution Reporting API](#)

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](#)
 PA API - [Protected Audience API](#) (formerly FLEDGE)
 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
Commitments	Section G of the Commitments are imperative to the viability of the Privacy Sandbox. Without a guarantee that Google's own ads business will operate exclusively on Sandbox technologies, the risk of ever-decreasing utility are raised as are the possibility of Google's divestment of the technology. Such a divestiture or reduction in utility would be an existential threat to privacy-forward addressability on the open web.	The Commitments do not guarantee that Google's own ads business will operate <i>exclusively</i> on Privacy Sandbox technologies. Google intends to use a portfolio approach to addressability, which will include the Privacy Sandbox technologies, the same way third parties can and do use. We understand a portfolio approach to be common across the ads ecosystem. We believe it remains important for developers to have privacy-preserving tools and technologies. We'll continue to make the Privacy Sandbox APIs available and invest in them to further improve privacy and utility.
Governance	The proposed governance model does not include specific mechanisms for accountability in formal	This is not correct. Both (i) the decision-making system and associated publications and (ii) the appeals process provide specific mechanisms for accountability. Furthermore, the Monitoring Trustee will oversee their functioning in detail.

	consultation or appeals processes.	
Governance	Feedback that the model does not contain provisions for the creation and maintenance of a cross-platform standard.	No governance model can compel other actors, in this case browsers, to adopt a standard. Thus we have not proposed a model that requires cross-platform adoption of standards. Google will continue to participate in standards forums where making proposals and sharing experience implementing proposals is a common activity in the process.
Governance	It is recommended to extend the consultation period to at least 2 months. The proposed governance model does not provide the ecosystem with adequate time to analyze the impacts of the proposed changes.	The three-week period is not the entire feedback period for a given change since existing feedback cycles (which are much longer) will continue. The consultation process instead offers a new, formal feedback window within the existing process for strategic decisions. As such, third parties will continue to be able to provide input through various forums (including GitHub, W3C and ads standards bodies like IAB Tech Lab) during the course of feature development. Structuring the feedback cycles in this way gives the ecosystem an opportunity to analyse and share their views on a proposed change without material delay to the development process.
Governance	Request for details regarding future governance plans.	A summary of the proposed governance model is set out in the CMA's Q2/Q3 2024 report (pages 3-5 here).
Exception Request	Request for an exception to access third-party cookies (3PCs) for their consented users.	Consenting to device access and storage or specific data processing purposes doesn't as such indicate a user wants to override their 3PC setting in Chrome. Allowing site-level override of a user's 3PC settings would create considerable potential for misuse, and it would be infeasible for Chrome to audit all sites' behavior that might lead to a request for exception.
Privacy Sandbox	Request for Privacy Sandbox API opt-in rates.	We have no plans to share this information with the ecosystem. Developers are welcome to call the APIs where they have code deployed today to estimate availability of the Privacy Sandbox APIs.

Origin Trial	Is there a plan to extend the origin trial?	The origin trial has been extended until April 14, 2025.
Privacy Sandbox	Request for a concise, non-technical explanation of Privacy Sandbox that highlights its business value and secures executive buy-in.	We have recently added a Business Resources section to privacysandbox.com here which provides this information.
Mode B	When a browser is in “Mode B”, will the current cookie jar (1PC, 3PC, local storage) stay or be wiped?	The current cookie jar will not be wiped. 3PCs will remain accessible in their first-party context.
Updated approach to 3PCs on Chrome	Will 3PCs be completely removed from Chrome in the future?	We are proposing an updated approach that elevates user choice. As set out here , instead of deprecating 3PCs, we would introduce a new experience in Chrome that lets people make an informed choice that applies across their web browsing, and they’d be able to adjust that choice at any time. We’re discussing this new path with regulators, and will engage with the industry ahead of rolling this out.
Chrome Testing	Request for continued availability of Chrome-facilitated Testing Labels.	The Privacy Sandbox team understands that companies would like to continue using the cookie deprecation labels . The process to extend the availability of the labels is similar to extending an origin trial. As part of this process, the experiment may only be extended for three Chrome milestones at a time . For example, Chrome’s most recent Intent to Extend Experiment (I2EE) for cookie deprecation labels was extended for Chrome M130-M132, inclusive. This enables support for the labels until the M133 stable release in early February. Additional extensions will run through the same process, so we recommend following along in the blink-dev email group for updates.

Enrollment & Attestation

No feedback received this quarter.

Show Relevant Content & Ads

Topics

Feedback Theme	Summary	Chrome Response
Specs	Is the classifier model shared between Android (by app name) and Chrome (by host name)?	No, they're separate models as they have different taxonomies.
Granularity of Topics taxonomy	Topics are too generic to be useful when leveraged with first-party information.	The Topics taxonomy seeks to balance utility and privacy. While we have evaluated possible mechanisms to make Topics more specific, we ultimately decided not to due to security and privacy considerations, among other concerns. Ad techs can unlock the best results by combining all available tools, such as machine learning and privacy-safe signals from privacy-preserving APIs, along with contextual data, creative data, and first-party data. Further guidance on this is available here.
API Usage	Topics API has low coverage.	Typical reasons for low coverage include: - User controls/opt-out - Publisher controls/opt-out - Site eligibility (the following sites aren't approved to match to Topics: insecure sites; WebView; Chrome on iOS, and Incognito mode) - User limitations (Chrome users who are under 18 or who are using Incognito mode cannot be observed and assigned Topics) - Seller observation requirement (the caller must have seen the user before on site associated with an eligible topic) - Implementation recency (allowing ~4 weeks to ramp up for caller observation to scale)
API Usage	Request for information on usage of Topics API as Networking tab seems to show there is a call sent and caught but chrome://topics-internals/ does not show observer recorded.	When using the HTTP header mechanism to interact with the Topics API, the topics are sent in the Sec-Browsing-Topics request header, but they are only observed if the server replies with the Observe-Browsing-Topics: ?1 response header. This is set out in further detail here .

Chromium Involvement	For Desktop, would Chrome share with other browsers based on Chromium the same observation and ranking context? For Mobile, would Android Chrome share with other Android browsers based on Chromium / In-App Chromium the same observation and ranking context?	Chrome does not share Topics data with other browsers on the device.
Specs	How does the Topics API decide if a page view by a user is regarded as 'topic history entry'?	To be eligible for the weekly topics calculation, a page visit must have an 'observe' call (can be from any caller). Without an 'observe' call, the visit won't be considered for topic history.
Security	How does the Topics API prevent one caller from getting other callers' observing topics?	We have provided an explanation here .
Taxonomy	How is the tree structure taxonomy within the Topics API used in the observation in each epoch?	When calculating the top 5 topics, it only considers the original topics provided by the classifier, and the rankings are determined by (i) the frequency of page visits, and (ii) a prioritization score (see the specification). When calculating the observed-by-callers of each of the top 5 topics, it includes callers who observed either the main topic or any of its descendant topics.
Specs	Request for additional information regarding the 5% random noise on the response.	There are always 5 top topics for each epoch. If the browsing history doesn't provide for 5 topics, then topics are chosen at random until there are 5. We call these padded topics. Callers will not receive one of these randomly padded topics unless they had observed (called the API on) the user on a site with the topic in the last few weeks. When the API is called, a per-user, per-site, per-epoch hash is calculated. If that hash is less than the probability of returning a noised topic, then the per-user, per-site, per-epoch random

		topic is selected to return. However, the noised topic is only returned (e.g., isn't filtered) if the caller had observed the corresponding un-noised topic for that user/site/epoch (see explainer). This filtering ensures that noised topics are returned 5% of the time for the given caller, regardless of its observational capability.
Specs	How does cross week duplicated topics work? Is the API picking independently among weeks and then merging?	Each week's (epoch's) topics are calculated independently. The particular topic chosen to be returned from each epoch depends on the site the caller is on. We do not take into account that the topic might be repeated across epochs for a given caller (and should therefore consider selecting a different topic) but we welcome additional feedback on this issue here.

Protected Audience API

Feedback Theme	Summary	Chrome Response
A/B Testing	The Shared Storage solution described here adds latency and has a high failure rate (i.e., a significant proportion of traffic ends up without a population ID). A low entropy ID (e.g., 3 bits) could be sufficient for effective A/B testing without significant privacy impact.	We believe the Shared Storage solution remains a viable approach, but we are considering this request and welcome additional feedback from the ecosystem if this feature is a high priority.
Reporting	Request for additional bits in reportWin(), particularly as the understanding is that the new click and display reporting in PA API will use 6-8 bits, effectively reducing the remaining bits available for other PA API reporting.	We are no longer considering increasing modelingSignals bits beyond the current 12 bits due to privacy concerns. We invite the ecosystem to provide feedback on our Private Model Training proposal, which aims at supporting ML training needs in a secure environment without a Privacy Sandbox-imposed limit on bits.
Interest Groups	Requesting 90 day life cycles for Interest Groups (IGs) as 30 days is not long enough.	As we've mentioned in our blog post , we plan to extend the IG lifespan to 90 days, and we have created an explainer here .

		We are currently working on the implementation, and will share a launch timeline when available.
Interest Groups	Requesting dynamic updates for IG delegation.	We are aware of this request from multiple stakeholders and are researching a solution. We will share updates on GitHub as this develops and welcome additional feedback from the ecosystem.
On-Device	Demonstrate more value for the ecosystem to continue investing in on-device PA API solutions.	The Privacy Sandbox team continues its development of Privacy-Enhancing Technology (PET) based APIs, including PA API, to offer developers more privacy-preserving options in the browser. Those technologies are generally available across Chrome browsers now on a large scale (not just 1% as some developers have misunderstood). Whether or not users enable 3PCs, developers may choose to incorporate PET-based solutions into their products, just as many companies are choosing to adopt other PET-based solutions outside of the browser. For many developers, they benefit already from investing in on-device PA API solutions by extending their deterministic first-party audience signal for improved reach across sites. We understand that some developers will only choose to use the Privacy Sandbox APIs or other PET-based solutions if more 3PCs are disabled, and these developers are waiting for information enabling them to speculate how many browsers will retain 3PCs. We recognize that those developers will wait until they find the information they are looking for in order to make product decisions.
Interest Groups	Allow SSPs to have any part in IG creation and the roles associated with them. SSPs see this as an important part of their value add, and would like the ability to help publishers sell IGs via their SSPs.	We have received the request to support more advanced IG delegations from multiple stakeholders, and we see the added value of SSPs contributing to this process. We are researching to find the best solution that allows different parties to participate in the audience extension process. We will share updates on GitHub as this develops and welcome additional feedback from the

		ecosystem.
Reporting	Discrepancies in the number of reports of "non-zero bids" between forDebuggingOnly and Private Aggregation API.	We expect a discrepancy to exist for two reasons: First, the Private Aggregation API debug mode is only available when there are 3PCs allowed on-device, while forDebuggingOnly API is always available unsampled (this last behavior will eventually change as detailed here). Second, Private Aggregation API has contribution limits while forDebuggingOnly doesn't. However, this feedback indicates there may be something else causing an unexpected discrepancy and we are working with a third-party stakeholder to resolve this issue.
Clickiness	As mentioned in the updated proposal to clickiness signal , views and clicks would be registered by returning a new HTTP response header to requests initiated by the "attributionsrc" attribute which are eligible", and this response header would include a list of origins that can be used to indicate which other parties can include these views and clicks in their aggregated counts. Does this mean that the ad tech can set the origins arbitrarily?	It is specified in the clickiness explainer that an ad tech that's contributing a view or click event to the aggregated view and click counts (a "providing origin") may include an optional parameter with the the response header that allows them to specify "one or more IG owner origins for which this event may be included in the computed view and click counts that will be provided to their generateBid() invocations in Protected Audience auctions" ("eligible origins"). That said, these views and clicks are not automatically included in the view and click counts. Rather, each ad tech must specify, in their IGs, the set of "providing origins" from which view and click events should be included, and only data from these providing origins contribute to the view and click counts presented to that ad tech's <code>generateBid()</code> calls. This mechanism requires an agreement on both sides, and prevents malicious players from influencing view and click counts for other adtechs. This also means that a "providing" ad tech that sets "eligible origins" arbitrarily won't have an impact on those eligible origins' view and click counts unless those eligible origins

		explicitly and intentionally include that providing origin in their IG.
Private Model Training	There are cases when DP-SGD (Differential Privacy - Stochastic Gradient Descent) would make the training process considerably slower as it destroys the sparsity of the gradients calculated during backprop. Are there any techniques that are being considered to navigate this or thoughts about this concern?	We are aware of some techniques to preserve sparsity in DP-SGD (e.g., this one), and we are exploring supporting these kinds of techniques in the private model training infrastructure. We will share updates as this develops and we welcome additional feedback here.
Negative Targeting	Request for updates on rolling out the IG negative filtering functionality described here .	As set out in our response here, we plan to support negative IGs in PA API bids. We will share a launch timeline when available and we welcome additional feedback here.
Bidding	Is it possible to combine multiple IGs when bidding?	This is not currently possible with PA API. PA API is based on the design that the IG relates to information a single site knows about the user, which has been core in discussions with the ecosystem at large. This approach allows ad techs to build a variety of solutions that help advertisers extend their 1P audiences across the web. We are aware that Microsoft's Ads Selection API proposal does propose a different design where audiences are based on information across sites. While we will continue to monitor their approach, we would want to see more discussion from the ecosystem, including the privacy community, before we consider similar changes to Chrome.
Native Ads	Concerns regarding whether or not PA API can adequately support the diverse formats and rendering requirements of Native Ads and whether	We are actively researching providing further support for Native Ads, and we welcome additional feedback from the ecosystem.

	PA API allows for the creative flexibility and optimization needed for effective native advertising.	
Reporting	Request to improve robustness of reporting scripts which compete for resources with bidding scripts and might be lost when the running auction frame is navigated away.	We hope to post a response to GitHub soon, but we don't envision these concerns significantly affecting valid reporting in practice
Macro Replacement	Auction config passed macro replacements not working with some 3P configurations.	The root cause was that not all Mode A/B labeled traffic had this feature turned on. We have recently decided to turn on this (and other features in the same situation) on all Mode A/B labeled traffic. We anticipate making this change during Q1 2025. We welcome additional feedback here .
Documentation	Requesting clarification as there appears to be a discrepancy in the documentation regarding the unit of measurement for the recency value in the browserSignals object within generateBid().	We have responded to this in further detail here and updated our documentation accordingly. The correct answer is that the unit of measurement is milliseconds.
Reporting	Request for 3P reporting; while DSPs and SSPs receive impression notifications from Chrome, middle-layer technical providers by default don't.	We are currently discussing this feature request here and welcome additional feedback.
Interest Groups	Request for guidance on how to exclude interestGroupBuyers dynamically when using parallel IG auction workflows.	We have provided guidance on this here .
Native Ads	Independent PA API auctions for a given page load prevents same-page ad filtering.	Further Native Ads support, including recommendation widgets known as "native" and that have multiple ads in one unit, is an active area of research and we acknowledge the

		current design may not support same-page ad filtering, and other protections like Fenced Frames may also prevent this further.
Native Ads	Existing PA API features like creative scanning, reporting, etc., that rely on URL-based signals may need to be adapted to handle JSON objects used in native ad creatives.	Further Native Ads support is an active area of research and we are assessing the feasibility of adapting PA API to aid native ad rendering.
Ad Verification	3P brand safety in PA API is affected due to latency and caching limitations of perBuyerSignals, which is problematic for dynamic content.	We recognize SSPs' and DSPs' need to determine an optimal TTL for caching that balances traffic shaping goals and brand safety max TTLs to ensure that cached data remains relevant for brand safety. We are exploring this issue and will share updates as it develops.
Ad Verification	FullpageURL macro replacement by SSPs is needed in order to conduct post-bid brand safety measurement.	deprecatedReplaceInURN is the current suggestion for SSPs to provide this signal.
Ad Verification	Lack of standardization in macro formats used by SSPs for post-bid verification may potentially cause inconsistencies and errors in data processing and analysis.	SSPs and Ad Verifiers should collaborate directly to define clear guidelines and specifications for macro usage to drive standardization of macro formats across SSPs to ensure consistency and prevent errors in data processing and analysis. This is an activity which ad standards organizations like IAB Tech Lab are well-suited for.
Ad Verification	Advertisers and Ad verifiers require a mechanism to link pre-bid and post-bid checks for the same publisher context for debugging and troubleshooting.	One option for post-bid verification is via auction- and campaign-based signals incorporated into event-level reporting. This may enable joins with pre-bid decision logs. We are exploring possible patterns for achieving this and will share updates as it develops.
Ad Verification	Request for exploring Trusted Key-Value (TKV) server solutions (DSP-owned and Ad Verifier-owned) for pre-bid and addressing fenced frames limitations for post-bid.	We are evaluating this request and welcome additional feedback from the ecosystem here to find a solution that could support pre-bid brand safety, and ease coordination requirements between parties.

Native Ads	Request for sell-side post-bid rendering audit for native ads.	Further Native Ads support is an active area of research and we are considering adapting existing features like this one to aid native ad rendering.
------------	--	--

Protected Auction (B&A Services)

Feedback Theme	Summary	Chrome Response
Latency	There has not been adequate mitigations to latency. While B&A Services may mitigate this problem in the long run, Google has not committed to its widespread availability prior to changes to 3PCs on Chrome.	We have made several improvements to on-device latency in the past few quarters. We are also building and scaling B&A services as necessary. We recently updated our latency best practices guide which includes more information on how to take advantage of these features. We are also continuing to develop new latency improvements, some of which can be seen here .
(Also reported in previous quarters) Auction Security	Request for further clarification on approaches to prevent/mitigate attempts to tamper with the on-device auction (e.g. including whether Google considers this a significant risk).	Our response is unchanged from previous quarters: “The reporting mechanisms of PA API ads retain the information used to distinguish humans from bot traffic today. Furthermore, current domain-based techniques of including or excluding domains can be used in PA API. This is described in more detail in our response to IAB Tech Lab's report on Privacy Sandbox .”
On-Premise Solutions	Concerns regarding how the largest suppliers might not adopt Privacy Sandbox or B&A due to the lack of support for private cloud, and the lengthy and opaque roadmap towards supporting it.	We are committed to expanding the infrastructures that Privacy Sandbox services run on; we have recently announced support for Azure cloud and continue our investigation into possible solutions for providing similar privacy and security safeguards for private clouds. Since our communication in October, we have made progress as we continue researching potential approaches to secure the privacy of Chrome users in an On-Premise Trusted Execution Environment (TEE). We now find ourselves at a place in our research where we want to validate different approaches with ecosystem stakeholders, and we plan to begin gathering input in Q1. Ecosystem feedback and collaboration will help refine any possible

		solutions.
Testing	Is it possible to create TEEs for testing PA API reporting solutions (Real Time Reporting and Private Aggregation)?	For Aggregation Service testing, ad techs can use test data and Local Testing tools to generate summary reports for functional testing. Please refer to the Local Testing Codelab here. For testing Aggregation in TEE, ad techs will need to complete the enrollment process as mentioned in the Codelab prerequisites above. We welcome additional feedback on this request here.
Deal K/V Integration	Request ability to pull deal based information from KV for business use cases.	We are evaluating this feature request and will provide an update on GitHub.
No-win Deal Measure	Requesting signal or ability to understand when a SSP didn't win and why.	We are evaluating this request and will provide an update on GitHub.
Feature Request	Request for Privacy Sandbox to provide a dictionary structure to help match a group of ads to the respective set of Deal IDs.	We are evaluating this request, together with other ways to reduce IG size with regards to storing deal IDs. We will provide an update on GitHub.
Performance	Solutions to measure possible missed ad auction opportunities, possibly due to bidding script size.	We are evaluating this request and welcome additional feedback here .
Specs	Currently B&A only reads the prevWins field instead of the latest field prevWinMs that replaced prevWins in the spec .	It is correct that B&A doesn't pass the duration in milliseconds in prevWins to generatebid(). Chrome sends the duration in seconds to ensure less overhead on payload, the fix here is for B&A to do the conversion from seconds to milliseconds. B&A would provide both prevWins and prevWinMs in browserSignals to make this compatible with on-device auctions. Note, even for on-device auctions for the web, prevWins and prevWinMs correspond to the same value and $\text{prevWinMs} = \text{prevWins} * 1000$. We are prioritizing a fix.
Documentation	The documentation is not	This codelab has been published as a guide to

	clear on how to test the Seller Front End (SFE), it would be helpful to have additional testing guidance right after completing the deployment as well as how to use Bazel for builds.	make it easier for the broader ecosystem to test their SFEs.
Deployment	Are there plans to provide pre-built binaries for B&A?	We are considering providing pre-built binaries for B&A, however we do not have concrete timelines for this. Until then, ad techs can build the binaries on their own and validate using the provided hashes.
Deployment	Should all orchestration types (server, client, mixed) be supported or is there one that should be prioritized over the others?	We do not have any specific recommendations on which modes the ad tech implements. The choice depends on various factors, but, ultimately, comes down to what's in the best interest of your customers.
Testing	Seeking clarification regarding failed tests during B&A build.	This could be a result of a flaky test. We have advised the ad tech to use the <code>--no-tests</code> flag to all <code>build_and_test_all_in_docker</code> build commands to skip running the tests.
Logs	Seeking clarification why the logs in log explorer on GCP are tagged to the VM instance running the SFE when in test mode but in prod mode the logs are not tagged to the VM instance.	It's hard to generalize as it depends on what exactly was seen but in general: - the logs from non_prod are probably the stderr logs redirected by the cloud provider (in this case, GCP) and GCP added the tag. - Logs produced by the VM are generally tagged with the VM instance, whereas logs produced by the binary are not tagged by GCP. - the logs from prod are exported by Open Telemetry in TEE, which have different tags. Here are some details of what is available at non_prod and prod.
B&A	403 error on missing secrets when OTEL logging is disabled.	This is now fixed as of 4.1 update and the documentation has been edited accordingly.
B&A	Missing outputs.tf file for GCP terraform configuration	This is now fixed.

	leads to error.	
Testing	Error when fetching private key in test mode.	In these instances please make sure servers are running with TEST_MODE=true. See explainer here .
Roadmap	Are there plans for getInterestGroupAdAuction Data to accept more than one seller origin and return a map of seller origin to B&A payload ciphertext?	Yes, adding support for allowing navigator.getInterestGroupAdAuctionData() to accept multiple sellers is planned.
KV Specs	Can KV URL (trustedScoringSignalsURL) be potentially delivered as a promise?	See explanation provided here .
B&A	Requesting creation of a new platform header to indicate to the B&A seller side what capabilities the client (browser) requires for a private auction to occur.	We are currently discussing this feature request here and welcome additional feedback.
Traffic Shaping	Proposal to drive down incremental costs from hosting B&A servers particularly for DSPs.	We are currently discussing this proposal here and welcome additional feedback.
Bring-Your-Own-Binary	Consider updating the explainer to explicitly address that all binaries are supported.	This is now updated, see the explainer here .
KV Calls	Does `generateBid()` wait for all Key-Value (KV) store calls to finish, or run independently? How does KV batching affect its timing?	See explanation provided here .
Performance	Request for additional documentation about re-using bidding scripts and recommending setting cache control headers on scripts.	Documentation has been added here .
Performance	Request to consider and	We are currently discussing this feature request

	explore the ability to load resources (e.g., bidding scripts) asynchronously in order to reduce on-device auction latency.	here and welcome additional feedback.
Consent Logging	Clarification needed for error seen when trying to use consent logging by setting the <code>CONSENTED_DEBUG_TOKEN</code> .	In these instances check that <code>CONSENTED_DEBUG_TOKEN</code> is present in the secret manager, <code>ENABLE_OTEL_BASED_LOGGING</code> set to true and <code>TELEMETRY_CONFIG</code> set to mode: PROD . See the explainer here which refers to the source here .
Logs	Are <code>forDebuggingOnly</code> events available through B&A?	<code>forDebuggingOnly</code> for B&A has been available for single seller auctions since April 2024. Our plan is to enable it for multiseller auctions very soon.
Worklet Logs	Request to make worklet logging compatible with ChromeDriver.	We are evaluating this request and welcome additional feedback here .

Measuring Digital Ads

Attribution Reporting (and other APIs)

Feedback Theme	Summary	Chrome Response
Debug Reports	How will ARA Debug reports be available to ad techs following the updated approach to 3PCs on Chrome? Should ad techs still have access to ARA Debug Reports for users who retain 3PCs and have the Privacy Sandbox APIs enabled?	Ad techs will have access to cookie-based and cookieless debugging solutions in respect of users with both 3PCs and ARA enabled. When cookies are off, they will only have access to the aggregate debug solution . Further details of the debug solutions are available here and here .
Feature Detection	Request for guidance on how to do feature detection for ARA on the server-side.	Currently ARA feature support can be identified based on using the Chrome version seen in the user agent string. We welcome additional ecosystem feedback

		regarding this.
Feature Request	Request for the <code>source_registration_time</code> used in ARA Aggregate <code>shared_info</code> to be more granular e.g. rounded down to one hour or one minute (as opposed to a one day) as well as configurable to take timezone into consideration (currently only UTC).	Rounding the <code>source_registration_time</code> field to the nearest day is a privacy mechanism used to mitigate the ability for an ad tech to be able to tie a specific user to a specific source registration. Currently the <code>source_registration_time</code> is based on UTC and an ad tech can adjust their ad reports to account for this. We welcome additional ecosystem feedback regarding this request here.
Specs	Request to clarify specification of "trigger_data" and "priority" especially when they are used with array value.	These fields don't accept array values. The square brackets in the explainer don't represent an array, but rather indicate that the text is not an example value, but a placeholder with a description. As the text in the square brackets suggests: - <code>trigger_data</code> is an int-64 - <code>priority</code> is a signed int-64 Neither of the fields accept array values. It's also worth considering using the header validator tool for ARA to experiment with different values and run into errors if the documentation is confusing.
Accelerated Mobile Pages (AMP) Ads Support	Does ARA support AMP ads?	Our proposal on how we could support AMP is available here and we welcome additional feedback.
Specs	Which URL will be considered as "source-site" for multi-layer embedded ads for ARA?	The URL from the top-level site.
(Also reported in previous quarters) Feature Request	Request for the <code>event_report_window</code> minimum value to be lowered from 3600 seconds (1 hour) to 1800 seconds (30 min).	Determining the minimum reporting window requires a balance of utility and privacy considerations. The minimum reporting window of 1 hour for event-level reports is essential to maintain privacy and mitigate against certain types of history reconstruction attacks.

		We welcome additional feedback on this request here .
Noise	Seeking a deeper understanding of how noise is implemented in ARA event-level reports to ensure accurate interpretation and utilization of the data.	Further details are available here and here .
Reporting	Aggregate Reports shared_info no longer contain source_registration_time by default.	This is due to API changes, and is set out in further detail here .
Reporting	The <code>filtering_id</code> is absent from the "Aggregatable Reports" tab of the chrome://attribution-internal s/ UI.	The <code>filtering_id</code> is currently visible in the "Trigger Registrations" tab details when you click on a row, which allows you to confirm its validity. We recognise the utility of showing it in the "Aggregatable Reports" tab, and have addressed this here .
Attribution Source	Request for clarification on how attribution source works.	An explanation is available here .
App to Web Attribution	Request for guidance for implementations where there is uncertainty whether sources will be OS or web.	Guidance is available here .

Aggregation Service

Feedback Theme	Summary	Chrome Response
Feature Request	Request for a configurable timeout for AgS and/or more visibility to job status for long-term runnings.	We are considering features to support monitoring long-running jobs. We welcome additional ecosystem feedback regarding this.
Terraform	Terraform trying to modify account IAM Policy even if service_account_token_creation_list is not set.	Developers can add their added permissions in their local modules/adtech_setup/main.tf file.

Documentation Request	Requesting documentation or a codelab explaining how to monitor aggregation service health.	A description of existing alarms that can be used to monitor the service and job health can be found in the relevant operator terraform files (alarms.tf and monitoring.tf) in the coordinator-services-and-shared-libraries repository. We will be publishing additional documentation and guidance on how to monitor aggregation service jobs.
Scaling	How to monitor the scaling issues?	We published an updated version of this guidance which documents the higher scale of the Aggregation Service. There is currently no direct signal that indicates a failure occurred because the machine cannot support the scale of the job. Indirect signals include: 90% memory consumption before a failure, a job that fails recurrently. We welcome additional ecosystem feedback regarding the need for such a signal.
Specs	What is the typical run time of AgS report batches?	Please refer to the guidance here .

Private Aggregation API

Feedback Theme	Summary	Chrome Response
Feature Request	Request to allow contributions of float values (including negative ones) to <code>contributeToHistogramOnEvent</code> with a sensitivity of 2^{16}	We are currently discussing this proposal here and welcome additional feedback.

Limit Covert Tracking

User Agent Reduction/User Agent Client Hints

No feedback received this quarter.

IP Protection (formerly Gnatcatcher)

Feedback Theme	Summary	Chrome Response
Geolocation	Request for IP Protection geolocation file.	The file mapping IPs to rough locations for IP Protection is available here . Please note that this file is updated periodically.

Bounce Tracking Mitigation

Feedback Theme	Summary	Chrome Response
Allow list	The updated position no longer addresses the allow list or similar mechanisms that would be independent of Google's decision-making process.	Google does not plan to have any allow lists associated with bounce tracking mitigations (BTM); the protections are applied uniformly to all domains.
Compliance	The ICO should have authority on privacy-related compliance.	Compliance status has no relation to the application of BTM and Google does not make any decisions regarding compliance in implementing BTM.
Competition	It appears that Google may be allowed to foreclose PET competitors using BTM (or other measures) and then exercise its discretion whether to allow them back into the market. The current appeals process may temporarily foreclose PET competitors from using BTM or similar measures.	The current BTM proposal is aimed at bounce tracking as a technique. While Google aims to avoid breaking certain use cases that may involve similar techniques, there is no plan for Google to provide individual exemptions from BTM. Thus the question of Google exercising discretion over the presence of competitors does not arise.

Strengthen cross-site privacy boundaries

Related Website Sets (formerly First-Party Sets)

Feedback Theme	Summary	Chrome Response
(Also reported in previous quarters)	The current limit of five associated domains is not high enough to achieve	Our response is similar to previous quarters: At present, we do not expect to increase the

Related Website Sets (RWS) Domain Limit	cross-site measurement use cases.	numeric limit. The limit was established based on user privacy considerations, feedback from ecosystem stakeholders in the W3C, and consideration of comparable implementations in other browsers. For additional information, please see our blog posts (1, 2). We recommend examining use cases that require cross-site cookie access beyond the numeric limit, and consider leveraging our guidance for identity use-cases, authenticated embeds, and advertising use cases. We welcome additional feedback on this issue here.
---	-----------------------------------	---

Fenced Frames API

Feedback Theme	Summary	Chrome Response
Native Ads	Native ad rendering in Fenced Frames poses challenges as inheriting the publisher's styling is limited due to limitations on communication between the iframe and the publisher's page.	Further support for native ads, including support following enforcement of Fenced Frames enforcement is an active area of research. We welcome additional feedback on this issue here.

Shared Storage API

Feedback Theme	Summary	Chrome Response
API Usage	The Shared Storage API is unavailable on some devices when other Privacy Sandbox APIs are functional.	This behavior is expected for a small subset of users (approximately 1%) who are part of the Shared Storage holdback experiment. This experimental setup is used to evaluate the performance and adoption of APIs in diverse scenarios.
API Usage	Do writes to Shared Storage occur under the publisher origin or the bidding script origin? Initial testing showed no writes when the publisher origin differed from the	This issue has been resolved, and only remains open in case of a possible devtools bug. Further details are available here. Shared Storage writes to the buyer origin in the bidding context of the <code>generateBid</code> call. The

	script origin.	write is not tied to the publisher origin, even if the publisher page resides on a different domain.
API Usage	What are the safeguards in place for a bad actor being able to read Shared Storage reports?	Shared Storage is partitioned by calling origin so a bad actor or ad tech cannot read Shared Storage data from another ad tech. Private Aggregation reports are sent directly to the same origin via TLS so they can't be intercepted.

CHIPS

Feedback Theme	Summary	Chrome Response
Partitioned Cookies	There is inconsistent handling of cookies across different localhost ports in Chrome and Firefox, particularly when using the Partitioned attribute.	Firefox treats localhost with different ports as distinct partition keys. While this behavior aligns with security principles; it deviates from the specification, and Chrome's approach. We expect to discuss this with other browsers in an HTML spec discussion , and will notify the ecosystem if this results in a change in the CHIPS partition key. We welcome additional feedback on this issue here .
Partitioned Cookies	The Clear-Site-Data draft incorrectly allows clearing beyond the partition of the emitting site, contradicting prior discussions referenced here .	This is a bug in the standards specification document, which we intend to fix soon. The correct behavior is specified in this section of the explainer, and aligns with the behavior aligned with other browsers on the storage partitioning explainer repository . Chrome's implementation already operates correctly.

FedCM

No feedback received this quarter.

Fight spam and fraud

Private State Token API (and other APIs)

No feedback received this quarter.

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 updates below:

Chrome-facilitated testing:

- On December 19, 2024, Google Marketing Platform published results from its standalone test using the Privacy Sandbox Attribution Reporting API on Floodlight Measurement for Non-O&O/Display. The results are available [here](#).

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 raised a number of concerns during the relevant period about impacts of the Privacy Sandbox changes, certain of which are captured in Google's responses to stakeholder feedback [above](#). Google is working with the CMA to resolve these concerns, following the process set out in paragraph 17(a)(ii) of the Commitments. The CMA has not notified Google of any concerns pursuant to paragraph 17(a)(iii) of the Commitments. The CMA has continued to raise detailed questions 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 to stakeholder concerns.

Updated approach to 3PCs on Chrome

As we announced in 2024, in light of feedback from stakeholders and the ecosystem, we have taken a decision to not deprecate 3PCs on Chrome. Instead we are proposing an updated approach that elevates user choice. We would introduce a new experience in Chrome that lets people make an informed choice that applies across their web browsing, and they'd be able to adjust that choice at any time.

Throughout the development of Privacy Sandbox we are engaging extensively with the CMA and ICO, in order to ensure that changes to Chrome continue to support competition and privacy in digital advertising. We are in discussions with the CMA on our updated approach and our Commitments to the CMA. We look forward to continuing to collaborate with the ecosystem on this project and we will share updates as they become available.

Status Meetings

The Commitments provide for Google and the CMA to schedule regular meetings at least once a month to discuss progress on the Privacy Sandbox proposals. In line with this requirement, Google and the CMA hold meetings to discuss a variety of topics relating to Privacy Sandbox

and Google's Commitments to the CMA, including technical, legal and procedural issues to assist the CMA in carrying out the regulatory scrutiny and oversight foreseen in the Commitments. Google and the CMA collaborate on the agendas for each meeting to ensure that adequate attention is given to each topic.

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 applicable at this time, 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 31 December 2024, 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 by: 
Signed.....
Full name
Date..... 

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