Privacy Sandbox Progress Report

Q2 & Q3 Reporting Period - April to September 2024

Prepared for the CMA, November 7, 2024

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 Q2 and Q3 2024 updates, covering the period from April 1 to September 30, 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 Q2 & Q3 2024 Timeline Updates	
April Timeline Updates	- Updated timeline to reflect new timing for Third-Party Cookie Phase Out in early 2025.
May Timeline Updates	- No changes.
June Timeline Updates	- No changes.
July Timeline Updates	- Removed the Third-Party Cookie Phase Out blue bars for Q4 2024 and beyond. - Added a text box for "Timeline update pending" which includes "Please read our July 2024 announcement for an important update regarding third-party cookies in Chrome" starting in Q4 2024.
August Timeline Updates	- No changes.
September Timeline Updates	- No changes.

Market Testing Grants

In an effort to encourage market participants to evaluate the Privacy Sandbox APIs, Google announced on July 18, 2023 that it would make 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, in line with the CMA's guidance to third parties on testing. As of the end of Q2 2024, grantees completed at least 8 consecutive weeks of testing between January 1 and May 31, 2024, and have submitted their results directly to the CMA. The Market Testing Grant program has now concluded.

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). As noted above, on July 22, 2024 Google announced that it would not deprecate third-party cookies (3PCs) on Chrome, and instead proposed to introduce an updated approach to elevate user choice. Therefore, with the agreement of the CMA, Google did not submit a public Q2 2024 report to the CMA in order to allow sufficient time for Google and the CMA to take into consideration the implications of Google's announcement.

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 API, Protected Audience API (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**
Third-party cookie deprecation (3PCD)	What are Google's plans for 3PCD and have the long term effects on the digital advertising industry been assessed?	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.
User Choice	The user choice announcement has impacted ecosystem interest in adopting Privacy Sandbox solutions. There is mixed feedback regarding the user choice announcement, including requests for features such as monitoring capabilities.	With the updated approach elevating user choice, it remains important for developers to have privacy-enhancing alternatives to cross-site identifiers. While we do not yet have details to share on what the new experience will look like, we do expect a significant increase of cookieless traffic in Chrome. This means the Privacy Sandbox APIs remain critical for businesses. We will continue to invest in Privacy Sandbox technologies to further improve privacy and utility.
User Choice UI	Questions about the timeline for the opt-out/consent features, the type of user option being considered, and how the UI will impact automated testing environments.	We do not have timeline updates to share at this time. Once we decided not to pursue 3PCD, we wanted to provide an update to the ecosystem as soon as possible. We'll share an update on the timeline for user choice as soon as we have one.
Chrome Testing	Request for continued availability of Chrome-facilitated Testing Labels to measure market adoption and economic impact of 3PCD post-H1 2024.	We are aware testers will want to continue using labeled browser groups for testing and coordination even as 2024 1H Chrome-facilitated testing has come to an end. We are evaluating next steps for labels in light of the user choice announcement. In the meantime, the Chrome team has published an intent to extend support for labeled browser groups through Chrome Milestone 132, which runs through January 2025.
Privacy Sandbox on Android	Privacy Sandbox on Android and Privacy Sandbox on Chrome are inextricably linked, and we cannot	*Please note that Privacy Sandbox on Android is not within the scope of Google's Commitments to the CMA.*

	properly assess Privacy Sandbox without Android. The typical customer journey, which involves cross-device and multi-touch aspects, is critical to both Privacy Sandbox on Chrome and Privacy Sandbox on Android.	If the feedback is specific to Android timelines and rollout, we have no updates to share at this time other than we continue to progress on Android, which we treat as an independent workstream for improving privacy. As we have previously noted, the availability of Privacy Sandbox APIs on Android will also be determined by the rate at which users update their devices, which is not in Google's control.
Mode B Traffic limited	Ad Auction Traffic available from Mode B has been lower than expected.	There are multiple reasons why auction volumes for the Protected Audience API (PA API) could be lower than expected, for example: - The companies that we know of who integrated PA API have only included banner formats. - Sell-side platforms may not always kick off an auction. - A browser may not have Interest Groups (IGs). - There may be no bids. However, we are unaware of anyone who attempted to test PA API and received no traffic.
Outage visibility	Visibility into outages and issues affecting the Privacy Sandbox APIs.	There is a public status page for the Privacy Sandbox APIs which have services outside the browser. The Chrome team places the highest priority on the reliability of our platform and all of the critical APIs used by major sites and services across the web, including the Privacy Sandbox technologies. Thus far there has only been one outage. It was related to a temporary configuration for testing at 1%. Soon the experimental configuration involved in this outage will no longer be needed, so we are confident this issue will not occur once the APIs are enabled in the normal manner in Chrome.
Cookie Graph Study	What is Chrome's perspective on the CookieGraph method as described in this paper within the Privacy Sandbox framework?	The paper raises some interesting points around the detection and prevalence of first-party (1P) cookies set by domains different from that visited by the user. As the paper points out, these cookies are extremely useful for gathering analytics and information of how users interact with a website. This data is crucial for

		developers to provide users with a better browsing experience. The main argument of the paper is flawed as it considers 1P cookies to be a cross-site tracking vector. However, this is only true under the very aggressive assumptions outlined on the paper: a) Users are willing to share their PII with the site. b) Devices have a stable fingerprint that can be used to track a user across sites. Note that these are vectors of re-identification that can be exploited without the use of 1P cookies (for example, through server-side data sharing), and need to be tackled separately from our current effort which is focused on state-based tracking mechanisms like 3PCs. Finally, we want to point out that the paper equates analytics and advertising cookies to tracking cookies and strictly necessary cookies as non-tracking cookies which may not necessarily be the case. Indeed, 1P analytics, or partitioned-to-site vendor services like store locator widgets, chat widgets, or load balancer cookies may often be limited to just one domain, and conversely some strictly necessary cookies might be cross-site tracking for anti-fraud purposes.
UX Changes	UX changes in Chrome 112 that places 1P cookie controls under the ‘on-device site data’ section of Chrome settings could make it more difficult for users to block all cookies.	This change was made as part of an effort to separate and elevate the controls for 3PCs (or unpartitioned storage) from all other kinds of site data. 3PC controls are elevated under the Privacy & Security panel; while controls for 1P cookies and all other kinds of site data - which critical site functionality typically depends on - are bundled under “On-device site data”. We will continue to monitor for feedback on this topic, but believe that the current separation strikes a good balance between discoverability of meaningful privacy controls, and preserving a functional browsing experience.

Billings and Payment	Billings and payments are not being fully tested as testers are more invested in testing other areas of Privacy Sandbox APIs.	When and what developers and companies choose to test is their choice. The APIs are generally available for testing and have been since September 2023.
Testing	Not all experimental traffic which DSPs are receiving from SSPs is labeled. Some DSPs have submitted that the share of experimental impressions which are unlabeled may be different across treatment and control groups.	Chrome cannot control whether companies forward labels in bid requests. We provide a method for getting a label from the browser. It is then up to the ecosystem to pass labels to partners if their partners cannot read those labels directly.
3PCD on Android WebView	Request for guidance on enabling the "Test Third Party Cookie Phaseout" flag in Android WebView for testing the deprecation.	3PCs are blocked by default in Android WebView.
Differential Privacy to mitigate risks in Model Training	Why is Differential Privacy used in Model Training?	Differential Privacy, combined with Trusted Execution Environments (TEEs), is essential in model training to prevent data leakage and secure sensitive information against threats. This approach ensures model weights cannot reveal individual user data.

Enrollment & Attestation

Feedback Theme	**Summary**	**Chrome Response**
Enrollment	Request for clarification on how attestation enrollment works between the origin that's enrolled versus the ad tech's origin with www subdomain.	The ad tech will only need to onboard on https://example.com. When they place their attestation in https://example.com/.well-known/privacy-sandbox-attestations.json, the https://www.example.com is covered since it's a subdomain.
API Spec	Suggestion to add a JSON schema for the attestation file to the repository.	We are evaluating this suggestion and welcome additional feedback here.

Show Relevant Content & Ads

Topics

Feedback Theme	Summary	Chrome Response
Topics Weighting	The most important thing to understand in Topics is the rarity of a given signal. The current design should evolve to allow adding a weight value beside each observed topic. The weight would be the relative weight of a given topic for a browser compared to all browsers using the topic.	We would like to understand further why the rarity of a signal is the most important signal. We welcome additional feedback from the ecosystem on the utility of this use case here.
Topics Reliability	Google needs to provide stronger assurances around the reliability of Topics over time.	Changes to our APIs will continue to be made based on ecosystem feedback and will be discussed publicly in advance of the changes. Our proposal for a revised governance structure would provide additional assurances.
Classifier	Publishers' sites are often misclassified or assigned Topics too high-level to serve any meaningful purposes.	As set out in our explainer on Topics classification, sites are classified through a combination of a human-curated override list, containing the most popular sites, and an on-device, machine learning model. Chrome continues to evaluate options for sites to contribute to Topics classification. Any utility improvements must be weighed against the privacy and abuse risks. For example, a few of the risks include: - sites using self-labeling as a method to encode different (and potentially sensitive) meanings into topics; and - sites attacking topics in order to blunt its usefulness for others (e.g., spamming the user's topics with meaningless noise). The public can inspect these components, with tooling available via a chrome://topics-internals or this colab. Through testing, we expect classification to improve over time, and we welcome feedback of examples of sites that may be miscategorized.

API Question	Concerns that Topics API gives persistent and anti-competitive benefits to SSPs that monetize bad content.	As with 3PCs, the browser is agnostic to whom it returns Topics to, as long as that entity is enrolled and attested.
(Also reported in previous quarters) Usefulness for different types of stakeholders	Because the Topics classifier currently uses only the page hostname to define the corresponding topics, large sites with diverse content are contributing generic topics while small sites are contributing niche topics with more advertising value.	Our response is similar to previous quarters: As 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.
Classifier	Smaller sites are frequently assigned an inaccurate classification or no classification so they cannot participate in the value exchange.	Regarding alleged harm, specific sites that are potentially misclassified are no more harmed by this than other sites, given that a site's contextual information will always be available for auctions on their site, which would provide comparable information to the correct topic, even in the case of misclassification. Topics are typically used to collect potentially useful advertising information from external websites, instead of their own sites.
Taxonomy Version	How can we access the taxonomy version to ensure backwards compatibility?	The taxonomy version is part of the request header sent with a topics-enabled fetch request. For example, if the header is "(1 2);v=chrome.1:2:5, ();p=P000000000" then the version is chrome.1:1:2. Where chrome.1 is the configuration version, the 2 is the taxonomy version, and 5 is the model version. This is described in the spec and has also been added to the explainer.
Topics Data	Request for clarification on how Topics data is updated.	The taxonomy update is not specified. This provides browser vendors with flexibility in implementation. Having said that, here are the heuristics of

		Chrome's taxonomy update from V1 to V2: - A single taxonomy tree is maintained for topics from both V1 and V2. - The same topic ID represents the same meaning. - The tree only grows – adding new topics or connections, never shrinking. - However, some topics or links could be "inactive" in a version, which could give the impression of deletion or reorganization. Examples: - "Pickup Trucks" now has "Trucks, Vans & SUVs" as an intermediate parent. - "Foreign Language Study" now has "Education" as a second parent, and its original parent "Reference" becomes inactive. Impact of "inactive" topics: These topics won't be used for new classification. However, they're still considered when enforcing user blocks: if a user blocked a topic in V1, its children will remain blocked in V2 (even if the child topic appears under a different parent in V2).
Classifier	Looking to understand causes and any corrective options available regarding erroneous classifications.	First, we'd like to point out that Chrome's determination of a site's topics is merely to be used as input to its Topics algorithm for determining a user's interests for advertising purposes. It is not developed for other, more general classification purposes. We're interested in the overall accuracy of our classification model, and try to improve their precision/recall where possible, but at the global level as opposed to the individual site classification level. This is because misclassification, when it does happen, does not harm the individual site that has been misclassified, rather it reduces the quality of the Topics signal when selecting an advertisement on other sites. When selecting ads on the misclassified site, the real topics of the site are already known to that site, and can be used as input to advertising queries. We welcome additional feedback here.

API Testing	Is Topics and in general the Privacy Sandbox APIs testable with Chromium?	The Topics API is not shipped with Chromium, it's shipped with Chrome.
Topics Caller	Request to improve added value of Topics leveraging TEE services for ad techs to support the cost of advanced analysis in privacy-compliant ways.	We have responded to similar feedback here. We considered inverse frequency, and ultimately after modeling inverse frequency we found that it did not correlate well with topic value as according to value provided by buyers and sellers. We welcome additional feedback here.
API Specs	Could browser interest-cohort setting block Topics API?	We have responded to this feedback here. Topics API is an evolution ofFLoC, and it honors FLoC's permissions policy. As set out in the explainer: "Note: The old Permissions-Policy: interest-cohort=() from FLoC will also forbid topic calculation."
Topics Ranking	When getting 'top 5 topics', would we count frequency of website visits based on each eligible caller, or always count based on the browser's whole visiting histories? Furthermore, are topics further ranked for each caller separately?	It's based on the frequency of a subset of browsing histories. A browsing history entry (a page) is eligible to participate only if the page had at least one Topics caller. Further detail on topics history storage is available here. As set out in our announcement on enhancements to the Topics API, the ranking depends on the frequency, and also on a binary priority level (see here and here for further detail). However, it does not depend on the frequency of callers. Please note that it doesn't mean all the callers can get all top 5 topics in the next epoch. As set out in the explainer "Only callers that observed the user visit a site about the topic in question within the past three weeks can receive the topic." The browser does need to track which caller observed which top 5 topics (corresponding to the top 5 topics with caller domains in the spec). We welcome additional feedback on this issue here.

Protected Audience API (formerly FLEDGE)

Feedback Theme	Summary	Chrome Response
(Also reported in previous quarters) Costs	More expensive to run TEEs in public clouds as opposed to on-premise ad tech data centers?	Our current TEE security model benefits from the practices of public cloud implementations (see more details in the public cloud TEE requirements explainer). For example, current hardware-based TEEs do not defend against all physical attacks. Our existing supported public cloud providers, AWS and GCP, designed and implemented mitigations for physical access risks, including from employees. While some ad techs have mentioned to us that running cloud services is more expensive than on-premise ad tech data centers, other ad techs run on public clouds whether it's for cost or other benefits. We continue to evaluate options for expanding our TEE support, including outside of public clouds. As part of that, we are researching on-premise data centers, and are engaging with the ecosystem to explore potential solutions for offering such support. At this current stage of research, we cannot guarantee that this will result in a workable solution for the ecosystem.
PA API & Google Ad Manager (GAM)	GAM is unable to achieve a fair market outcome. GAM fails to serve ads in a timely manner, report how many ads it served using PA API, and does not offer configurability as to which method it will select to serve an ad, e.g. by turning off PA API for certain slots.	**Google Ad Manager Response:** GAM has and continues to work on optimizing latency when serving ads via the PA API so that the additional publisher revenue gain from PA API demand outweighs any costs incurred due to the additional PA API auction latency. Our initial testing does indicate that publishers see a net revenue benefit from PA API on traffic without 3PCs, indicating the additional demand from PA API outweighs any costs due to latency. Further details on our approach can be found in our FAQ. Additionally, GAM provides publishers reporting on ads served via the PA API. This includes ads served when GAM is a component seller, and ads served via other component sellers when

		GAM is a top-level seller. Further details on reporting can be found in our Help Center. Finally, GAM does allow publishers to enable or disable its use of the PA API on their traffic via an in-UI control (see our Help Center for details). We are open to considering feedback on further controls publishers may desire and will prioritize any feature requests in line with our standard feature prioritization process.
PA API & GAM/AdX	It appears Google has taken the position it will simply not buy any PA API impressions that GAM does not make the final selling decision on, much like AdWords only buys from the house. This seems purely an abuse of market position, as GAM/AdX could submit a component auction configuration to an alternative top-level seller like any other exchange.	**Google Ads Platform's Manager Response:** That is not Google's position. Google's buyside platforms (Google Ads and DV360) do buy impressions from non-Google exchanges. This is true for both PA API impressions and non-PA API impressions.
Market Response	Mozilla's concerns: Google's Protected Audience Protects Advertisers (and Google) More Than It Protects You.	We appreciate Mozilla's assessment and will continue to engage with Mozilla's feedback in public standards forums. A theme of their assessment is that the current implementation of PA API does not yet enforce all of the planned protections. Our go-to-market approach with PA API has sought to strike the right balance between encouraging adoption and implementing privacy protections as soon as practical. As part of this, we've established a roadmap for imposing privacy restrictions over time, in order to better facilitate integrations with the APIs as well as to give us time to collect more feedback we can incorporate into the future protections (e.g. VAST features in Fenced Frames). We also welcome Mozilla's more recent communications about its own approach to privacy and digital advertising: "A free and open internet shouldn't come at the expense of privacy" and "Improving online advertising through product and infrastructure".

(Also reported in previous quarters) Auction Results	Request for single auction to return multiple render URLs with their corresponding score, making it easier for native advertising to deduplicate and avoid UX and latency issues.	Our response is similar to previous quarters: Sharing multiple renderURLs, and their respective score, from a single PA API auction is something we considered but did not implement due to privacy concerns. We do understand the desire to avoid showing the same ad multiple times to a user on a single page and are evaluating this request. We welcome additional feedback from the ecosystem here on what additional support is needed in PA API to optimally support Native Advertising use case.
Performance	Concerns around latency impacting PA API.	We have heard concerns about latency and this is part of the reason that we have developed a number of features as part of the PA API which will make it possible for SSPs to both set limits on DSP latency as well as make improvements which can decrease latency. 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.
Top-level sellers	Google should enable publishers to choose alternative top-level PA API auction sellers.	PA API is agnostic to who initiates an auction both in single seller and multi-seller designs. Individual companies' choices about whether and how to support PA API auctions are their own.
(Also reported in previous quarters) Negative Targeting	Request for a solution to a use case where an advertiser does not want to display ads to a certain audience.	We support negative IG targeting through additional (contextual) bids, which solves the needs where an advertiser doesn't want to display ads to a certain audience. The details can be found in this explainer and this GitHub issue. We are also exploring solutions to support negative IG targeting for PA API bids, and welcome additional feedback here.
(Also reported in previous quarters) Native Advertising	Request for Fenced Frame support for Native Advertising.	We are considering supporting this use case and are discussing possible workarounds and solutions here.

WebView	Seeking clarification on the scenario where IG joined at Chrome was not available for Auction executed on WebView.	We do want to support these use cases once sufficient privacy infrastructure is available. We don't have any further announcement to make at this time but we welcome additional feedback here.
Negative IGs	Request to updateURL processing to support negative IGs via the nascent header feature.	We are evaluating this request and welcome additional feedback here.
Diversity Filtering	Request for guidance on how to implement diversity filtering when running native advertising in PA API with multi sellers and multi auctions.	We are discussing this request here and welcome additional feedback.
(Also reported in previous quarters) Blocking Filters	Request for guidance on how to enforce 'publisher blocking' (filters) rules when running native advertising in PA API with multi seller.	We have shared guidance here and welcome additional feedback.
Restrictions	Request to allow restrictions at the domain level rather than at the subdomain level.	Restrictions at the subdomain or origin level follow the basic security model of the web so that's our default design. We have discussed this in further detail here and here.
Trusted Bidding	Request for user agent and related client hints in trusted bidding signal requests.	We are tracking this feature request and welcome additional feedback here.
(Also reported in previous quarters) Multiple IGs	Use multiple IGs in the same bid.	This is not supported in PA API today, as it would result in a change to the underlying privacy model. We welcome additional discussion here.
(Also reported in previous quarters) Performance	Moving more logic to the client can harm the page performance and UX, potentially hurting the website SEO scores.	We are discussing this issue and welcome additional feedback here.

Auction Dynamics	PA API reduces information on auction dynamics (e.g. who participates, who wins each component auctioned etc.) which reduces traceability publishers and makes it hard to know whether deals are being kept.	We proposed a solution to the tracking of deals here. We plan to modify some existing fields and create some new fields within the IG object to store DealID and SeatIDs, and allow them to propagate from generateBid to scoreAd or egress via event-level reporting. This should provide adequate support for the deal's use case. We welcome feedback on other metadata that ad techs consider critical to auction dynamics and to keep having this traceability for publishers. We encourage ad techs to provide specific examples of metadata they are referring to and from which party to which party it needs to flow.
GAM	Concerns over the requirement to use GAM as the publisher ad server in order to access AdX demand.	**Response provided by Google Ad Manager:** GAM does not require that publishers use its ad server functionality in order to access its exchange functionality.
(Also reported in previous quarters) Component Auction	PA API component auction winners will be visible to GAM, raising concerns about unequal access to information.	Our response remains unchanged from previous quarters: **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 API auctions, we intend to keep our promise 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 any component auction, including our own, as explained in this update. Moreover, we do not use information about component auction configurations, including signals provided by buyers to SSPs, as part of

		our own auction. In fact, we would welcome changes to the PA API that allow component sellers to specify their component auction configurations in a way that is obfuscated from the top level seller.”
GAM	Will GAM request revenue share for running/reporting of top-level auctions if GAM has not participated in either the creation of IG or PA API component auction?	**Response provided by Google Ad Manager:** When publishers choose to use GAM as their ad server, GAM will run the top-level auction and charge an ad serving fee for its ad server functionality (the same ad serving fee that applies outside of PA API auctions). However, if the winning ad comes from a non-GAM component seller - meaning GAM has not participated in either the creation of IG or PA API component auction - GAM does not handle billing and does not charge a percent media fee.
Clickiness	Is the registration of the click events subject to the same differential privacy?	This feature is currently not planned to be subject to "k-anon" restrictions, because the "count of clicks" will only be available as a browserSignal inside the generateBid() function; it is not available as a new attribute in event-level reporting.
Performance	High egress costs due to unconditional response to contextual bid requests.	We cannot directly provide information on which DSPs have IGs due to privacy concerns. However, we are exploring alternative solutions that could provide insights while preserving privacy.
Native & Outstream Ads	Request for updates on Chrome’s perspective regarding native & outstream ads.	Chrome’s position is dependent on the use case in question. On Video, Chrome’s position is that our job is to encourage the ecosystem to converge on viable instream Video solutions using our APIs. Thus far, the only concrete proposal we are aware of is GAM’s proposal. On Native, we are actively collecting feedback here and plan to engage ad techs with more discovery steps soon.

Real-Time Monitoring (RTM)	Labeled traffic does not send RTM reports.	We are aware of this gap and are working to provide a solution. We will share an update when available here.
Audience Extension Support	Request for update on support for audience extension/seller-curated audiences in PA API.	We are working to provide a solution to this use case. We are collecting feedback from the ecosystem on what we should build and support. We will share an update when available and we welcome additional feedback here.
Debugging	In Chrome's developer tool, there is no panel to monitor the detailed performance of PA API. For example, the overall performance might be affected by the number of IGs or number of buyers.	While this specific feedback relates to the Chrome Developer Tool UI's capabilities to assist with monitoring, on October 7 we introduced the ability for ad techs to configure custom events that can be used as the basis for monitoring and alerting. Further details are available here and we hope that this update addresses a material portion of this feedback. We welcome any further feedback on this feature, whether related to the supported data points or the developer experience in the corresponding GitHub discussion here.
Signals	Concerns regarding whether or not DSPs can ensure perBuyerSignals is sent to SSPs independent of contextual auction results.	The contextual auction is assumed to have only one winning bid from one DSP, or better said one bid to try to beat with a subsequent PA API auction. For the PA API flow the SSP decides to invite any and all DSPs that they wish to see if they have demand (in the form of an on-device IG) to submit. It's entirely possible and actually very likely that a DSP that lost the contextual auction is "re-invited" to participate in the PA API auction. In this "re-invitation" is when the DSP, if it decides to accept, would forward on to the SSP any signals the Ad Verifier would want to make sure the DSP considers, if any exist for that campaign. In other words, in the PA API auction there is always a way for the DSP to submit perBuyerSignals to the SSP regardless of what transpired in the contextual auction.
Signals	Request to add prevClicks to browserSignals object	This request can be resolved by our proposal to support clickiness signals. We announced this

	passed to generatedBid().	feature in our recent blog post and corresponding explainer. We welcome additional feedback on this proposal here.
(Also reported in previous quarters) Modeling Signals	Request to increase the number of bits of modeling signals from 12 bits to 30 bits.	We have responded to this feedback with a counter proposal here. We are actively engaging with the industry to understand their views on our proposal, and are currently weighing the benefits to the industry against the impact on Chrome users and other stakeholders.
Documentation	Request for guidance on how to use Key/Value (K/V) Servers and TEEs.	Guidance on the use of TEEs in the context of K/V is available in the K/V service trust model design details here.
Lifetime of negative IGs	Request to extend lifetime of negative IGs to 365 days.	Negative IGs are used to prevent showing ads, but bad actors can still use it to reveal information about users, resulting in re-identification risks (e.g. one way for bad actors to attack is to just place high bids with negative IGs in them repeatedly to learn if a user has or hasn't visited certain sites). If we keep a 365-day TTL, then bad actors will have a lot more data about negative IGs which results in significantly bigger privacy risks. Therefore, we cannot support this request in order to protect user privacy.
API Specs	What can be inserted as values to be passed as part of perBuyerSignals? Can this be arbitrarily defined by the seller?	perBuyerSignals is the place for sellers to provide to buyers whatever information they want to make available inside the auction. It is for the ecosystem to decide what they wish to insert there, but we welcome additional discussion here.
Ad Size Macro Replacements	Seeking guidance around ad size macro replacements not working.	We will be sharing more details publicly soon.
Post Bid SSP Macro Replacement: Spoofing Top Level URL	What mechanisms can Chrome introduce to allow verification vendors to verify the top-level URL within the Privacy Sandbox framework?	We are currently discussing this welcome additional feedback here.

	Are there alternative data points or signals that can be used within Fenced Frames to ensure the accuracy of the SSP-provided top-level URL?	
Feature Request	Request to provide low-entropy UACH on updateURL fetches and on Real-Time Reporting postbacks.	These requests are under discussion here, and we welcome additional feedback here and here.
Feature Request	Request to have the trusted server consented debugging design to be activated when a given client has been triggered to send downsampled forDebuggingOnly event-level reports via forDebuggingOnly.reportAdAuctionWin() and forDebuggingOnly.reportAdAuctionLoss().	This is an active request we're currently tracking and will provide an update to the ecosystem when available. We welcome additional feedback here.
API Usage	Request for guidance on how to count unique user reach and impression reach.	We are discussing a proposal to address how to read IGs from within a shared storage worklet, which you could then send private aggregate reports on. Further details are available here and we welcome feedback on the proposal and its usefulness to the ecosystem.
API Usage	Lack of clarity on what publishers should test, which APIs are important, which one should be turned on and what is to come.	There are efforts underway to better support publishers and their roles in the ecosystem.
API Usage	Is it possible to add event listeners to Ad Auction Worklet events?	This is not possible via normal events but Chrome Devtools Protocol events will partially address this use case. Note that regular events are likely to impact isolation/privacy properties, but details are available here.

K-Anonymity	Seeking clarification on ad rendering requirements (e.g. at least 50 people would have seen the ad, if it were allowed to show).	The developer documentation provides an overview of our expectations for future developments. In particular it explains that the initial k-anonymity threshold is k=10 people. The blink-dev mailing list provides updates on what is happening live in Chrome. As set out in the k-anonymity mailing list thread, we are currently experimentally enforcing the k-anonymity requirement on about 1% of Chrome stable traffic, and never enforcing it on the explicitly labeled ("Mode A" and "Mode B") slices.
Chaffing	Can the TEE K/V chaffing be temporarily removed or reduced from having to call all N shards, to some amount that balances privacy protection against utility (i.e., K/V performance/cost)?	These types of requests are handled for only non-production instances where chaffing can be controlled. For production instances chaffing is still required. We can evaluate the situation once we receive feedback from non-production usage.
Chaffing	Request to add flag to disable chaffing from debug/non-prod K/V binary.	This flag is provided with the release 1.0.0.
API Bug	API stopped working after upgrading to Chrome 126, even though the API was enabled in settings.	This issue was found to be related to the "enable-fenced-frames" Chrome flag, which was enabled by users for development purposes. Resetting this flag to default will resolve the issue.
Reporting	Request to make real-time reporting API opt-in not seller-dependent for buyers.	This request is being considered here. The RTM solution was released recently and we welcome feedback in particular from those ad techs that have already onboarded to the feature.
Reporting	Request for 3P reporting; while DSPs and SSPs receive impression notifications from Chrome, middle-layer technical providers by default don't.	We are discussing this request and welcome additional feedback here.

Protected Auction Services

Feedback Theme	Summary	Chrome Response
TEEs	Google's requirement for manual onboarding under technical standards is a strong restriction on the choice of cloud vendor. The technical standards applied can be followed without a visit to the Bureau of Cloud Providers as Google seems to have in mind. The late delay of alternative providers in 2025 (earliest) is unacceptable because it will introduce network effects encouraging tipping to Google's solutions.	Aggregation Service is the only service required to run in a TEE service to address some ad-tech use cases. Aggregation Service supports both Amazon Web Services (AWS) and Google Cloud Platform (GCP). Based on feedback from ad techs, we believe such support is adequate at this stage. On additional cloud providers - Google published detailed criteria for TEEs on Public Cloud Providers. These are aimed to ensure that the TEE solution meets privacy and security goals of Privacy Sandbox. Specifically, Privacy Sandbox TEE servers process unencrypted cross-site user data (e.g. data from the publisher and advertiser sites for Aggregation Service). These need to be secure in order to meet the user privacy goals of the APIs. A secure environment is likewise necessary to ensure the APIs continue to protect companies' confidential business information (for example, preventing other PA API auction participants from accessing a buyer's proprietary business data). To the best of our knowledge, there is currently no TEE technology which fully protects user data from a potentially adversarial operator. Therefore, we include multiple requirements to validate the trustworthiness of the cloud provider. We are uncertain what "Bureau of Cloud Providers" refers to, and it is not part of the requirements. We welcome any feedback on the requirements. We also continue to evaluate support for new providers, including based on requests for submitted using the process defined in the explainer. So far, we have only received a request to support Azure, which we are evaluating.
B&A	It is difficult to assess the technical complexity and	To address these concerns, we have provided detailed explainers on GitHub explaining the

	feasibility of the B&A service as the design is continuing to evolve.	design of B&A, published timelines of availability and a roadmap of features supporting PA API. We are supporting ad techs who seek to deploy B&A and collecting feedback from the ecosystem on GitHub.
B&A	Looking for guidance and a better way to calculate the cost of using TEE for B&A in order to start using it or migrate to using it from on-device.	We recently published the K/V Server Instance Sizing Guide, which includes tooling to more accurately measure costs.
K/V Server	Ad-verifier requesting to be able to use the full page URL to the K/V server to perform ad verification.	We are currently evaluating the possibility of providing the full page URL to K/V server running in a TEE for ad verification purposes. The full page URL won't be available in K/V BYOS.
Auction Security	Seeking auction security features to ensure bad actors don't access sensitive data or act as impersonators - which features protect the auction from replay attacks and how can security safeguards be implemented?	B&A's current security model can protect auction integrity. B&A runs in a TEE that protects the confidentiality of ad techs' signals and code from malicious actors. In B&A architecture, an encrypted B&A request payload (request ciphertext) flows from the client through an untrusted ad server to SellerFrontEnd service (SFE, run by SSPs in TEE). The request ciphertext contains a timestamp based generation id. The SFE will examine the timestamp of the request and reject any replayed requests not within +/- n seconds of the server time. In addition to that, B&A can return a padded fixed size response payload for server to server communication. These solutions can help mitigate replay attacks through the system when a malicious actor tries to replay the same request payload to learn more about its contents. B&A is in the process of documenting and updating security models in explainers.
Signals via K/V Server	Request to include perBuyerSignals sent via K/V service as part of the trusted bidding signals request from Chrome.	We are evaluating the feasibility of including information from perBuyerSignals, transferred to K/V server running in a TEE including full page URL.

K/V Server	Request for a more phased adoption timeline for privacy constraints in K/V and B&A.	We understand the desire for a more phased approach to TKV adoption and appreciate your specific requests regarding K/V and B&A. However, after careful evaluation, we've decided not to relax the privacy protections in these APIs at this time. We believe these measures, such as chaffing, are crucial for safeguarding user privacy and maintaining trust in the Privacy Sandbox.
K/V Server	Seeking guidance on how to scale the K/V store through a compatible configuration.	The recently published K/V Server Instance Sizing Guide can help here. The tool will provide the QPS (noted as "RPS" in the explainer) at each combination of parameters.
K/V Server	Add Seller Information on the K/V Server request.	We are discussing this and welcome additional feedback here.
K/V + B&A Services	Request to clarify the release timeline or roadmap for K/V and B&A.	For both K/V and B&A, we have stages and timelines: For K/V Server in conjunction with on-device PA API auctions (vs B&A) the public timeline is available here. In terms of how "General Availability" is defined for K/V, see the Roadmap section which defines the feature set for Beta and GA. For B&A see the public timeline here and the roadmap here. We define Scale Testing as "full stable, production scale testing" -- see here for the specific feature set at this stage. B&A also has Alpha and Beta stages, so the Scale Testing will include the super-set of features of prior stages. For both K/V and B&A, let us know if these stage definitions help provide clarity as to what would be available when. If there are still gaps, please let us know. We are happy to make these more specific to help inform planning.

Measuring Digital Ads

Attribution Reporting (and other APIs)

Feedback Theme	Summary	Chrome Response
Market Response	The requirement for competing attribution systems to use only event-level reporting and summary/aggregate reporting within tight bounds is an arbitrary restriction on competition. It prevents real time device-specific retargeting and attribution at the event level, even if safeguards are in place to ensure data protection compliance (e.g. de-identification).	The design mentioned is based on the privacy goals of the API - e.g. protecting cross-site information being passed from one site to another. For example, ARA supports event-level attribution via event reports. Event reports are delayed at a minimum of one hour, which is necessary to make it difficult to associate the event-level report with the user's identity on the advertiser's site, using timing side channel attacks, as documented here. Additionally, there are other ways to do attribution, beyond ARA, such as directly collecting information from users who knowingly provide identifying data. We are open to feedback on use cases that cannot be achieved with the current privacy bounds of the Privacy Sandbox APIs.
Multi-Surface	Request for confirmation on whether or not ARA and Shared Storage APIs support multi-surface use cases and where this is evidenced.	Currently ARA and Shared Storage do not support multi-surface (cross device) attribution. Cross app and web attribution on the same device (via ARA) is supported.
(Also reported in previous quarters) Cross-Device	Does ARA support cross-device conversions?	Our response is similar to previous quarters: 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 currently have a timeline for cross-device support.
Scaling	Concerns about whether the Attribution Report API (ARA) is currently configured and can be reliably rolled out and scaled to service all Chrome users.	ARA is currently available to all Chrome users and running as expected. Additionally, we continue to test and monitor its reliability and scalability, as the number of ad tech companies testing ARA continues to increase. We welcome additional ecosystem feedback

		regarding this here.
(Also reported in previous quarters) Deduplication	Concerns on how ARA proposes to restrict the attribution mechanism on devices such that publishers are not able to effectively perform post-attribution logic for summary reports, including deduping multiple same-type conversions for the same ad click.	Our response remains unchanged from previous quarters: Deduplicating across devices and measurement pipelines is a known and current challenge that ad techs also face today with 3PCs. With ARA, 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 welcome additional ecosystem feedback regarding this here.
Conversion Tracking	Request for ability to operate with conversions from multiple domains.	We are discussing this request here and welcome additional ecosystem feedback.
Reporting	The browser waits at least two days but up to 30 days to send the conversion which can be cause for concern given the majority of the stakeholder advertisers are performance advertisers, which work in near real-time times.	The default settings for event-level reports have the following default reporting windows: 2 days, 7 days, and 30 days. With flexible event-level reporting ad techs can change the number and length of reporting windows from the default values. Reporting windows can be changed to a minimum of 1 hour which may help with performance advertiser use cases. We welcome additional ecosystem feedback regarding this here.
Online-to-Offline Attribution	Will there be any implementation options for online-to-offline advertising in ARA, or are there any other suggestions for measuring offline-to-online attribution?	Currently there are no plans to support online-to-offline measurement use cases in ARA. There are significant privacy and security challenges that need to be considered for this type of support. We welcome additional ecosystem feedback regarding use cases for this support here.
Debug Reporting	How to store and/or retrieve AdID in such a way that it's available to be accessed for Chrome (source/trigger)	In order to enable the debug reports, the ad tech must prove to us that they can already join the user across app and web, and this is done to ensure no new information is revealed by the

	registrations for app-to-web attribution?	debug reports. The ad tech can prove the join by providing a join key that is unique per user. This join key can be the AdID or can be a 1P join key. If the ad tech uses the AdID, Chrome does not natively support accessing the AdID from the browser and the API expects each ad tech to have their own method of passing the AdID during the web registration.
Bucket Granularity	Is it possible to use different bucket strategies per advertiser?	We recommend experimenting with different contribution budget scaling approaches to find the one that works best for your use cases. ARA was made with the intention of being flexible and customizable to satisfy a variety of ad-tech use cases. Therefore we recommend experimenting with different bucketing strategies per advertiser or per vertical. Using different bucketing strategies can be useful when advertisers have differences in measurement values they are interested in tracking and the volume of the measurement values.
Documentation	Request for additional documentation for implementing app<>web for ARA.	We have released documentation on App<>Web for ARA here.
Performance	The number of ARA-related requests can potentially be a heavy load on a publisher's server(s) relative to the number of keepalive requests that are necessary to power said site. Batching source events in a single request can help reduce load on a server. One potential idea is to allow an array of JSON-encoded objects	Batching source events based on specific logic is possible to a certain extent using JavaScript logic in combination with the API. We are currently evaluating this request and welcome additional feedback from the ecosystem here.
Feature Request	Suggestion for a workaround proposal due to no server-to-server integration support.	Currently we do not plan to implement support for server-to-server integration in ARA. There are many privacy challenges that need to be further considered to allow supporting server-to-server integration. We welcome feedback from the ecosystem regarding additional use cases for

		server-to-server support here.
Documentation	Request for a "quick-start" guide that explains the key parts of ARA/how to get up and running with a couple of simple use cases.	A quick-start guide for ARA is available here. We are working on improving this documentation further this year, and welcome additional feedback on specific use cases or scenarios that require additional documentation here.
API Usage	Request for recommendations on scaling contributions for many small values.	We recommend experimenting with different contribution budget scaling approaches to find the one that works best for your use cases. For the scenario of many small values we recommend experimenting with different values of epsilon to identify inflection points at which the noise from epsilon may be acceptable for your use case. Further details are available in our research paper on Summary report optimization in ARA.
Flexible Event-Level	When will Flexible Event-Level (multiple trigger specs) be implemented?	Currently Flexible Event-Level supports customizing the following registration configuration aspects: the number of reports that can be generated per source, the number of and length of reporting windows, and the cardinality of the trigger data. We are currently gathering additional ecosystem feedback regarding additional flexible event-level enhancements, but do not plan to implement any currently. We welcome additional feedback on use cases that might benefit from some of the flexible event-level enhancements listed here.
Bucket Processing	Request to consider capping aggregated contributions for buckets as well as future extensibility and backwards compatibility.	We are discussing this request and welcome additional feedback here.
Epsilon	What happens to the epsilon range once ARA changes to general availability?	ARA reached general availability on Chrome in Q3 2023. At this time, there is no plan to modify the epsilon value window. Our proposal for a revised governance structure would provide additional assurances where any modifications

		are envisaged.
Reporting	Trigger-unknown-error reports don't contain source attributes in the report body.	As set out in the specification, there is no requirement for other fields to be present in the report body for *trigger-unknown-error*. We recognise that the table describing fields in the report body was potentially misleading, as the source-related fields may or may not be present depending on the underlying cause of the error. For example, an internal error could occur before the source-matching logic happens, which would mean that no source data is available to populate the debug report with. The documentation has now been updated to clarify this.
API Usage	When working with two measurement goals, count and value, is the indication to split both contribution budget and epsilon?	When working with two measurement goals, we recommend splitting the contribution budget between them.
Reporting	Does ARA support multi-touch attribution or assist reports (a.k.a. contribution reporting)?	ARA does not currently support multi-touch attribution or assist reports. Currently we have no plans to implement this. We welcome additional feedback on use cases and their priority here.
API Bug	For ARA, the documentation states that XOR is used to combine source and trigger aggregation key-pieces, but in practice, OR is being used. This discrepancy leads to confusion and potential errors in implementation.	The documentation has been updated to reflect that this is an error.

Aggregation Service

Feedback Theme	**Summary**	**Chrome Response**
Aggregation Jobs	Request to allow multiple prefixes in aggregation jobs.	We are investigating this feedback and have posted a proposal here. We welcome feedback on the proposal here.
Feature Request	Request to change terraform	We are investigating this issue here and

	script so that if service_account_token_creator_list is not set (or is empty), then modification of account IAM policy is skipped.	welcome additional ecosystem feedback.
API Usage	Clarification needed regarding Terraform plan always showing changes.	This issue has been fixed in the AgS 2.8 release.
API Usage	Seeking recommendations for setting up per-advertiser settings for aggregation frequency with flexible contribution filtering.	Batching by advertiser is currently possible with ARA. Flexible contribution filtering could be used for more advanced cases where an ad tech wants to further segment contributions of a report by different frequencies. Ad techs can learn more about batching here and using filtering IDs with ARA here. We are also working on more documentation for filtering IDs.
Feature Request	Request support for `log_sum_exp` in Aggregation Service (AgS).	We have reached to this stakeholder for more details on the use case. We will provide an update once we have more details.
Feature Request	Request to be able to see more logs/insights/other metrics whenever there are issues on AgS and potential issues on a deployed AgS.	We have published updates to our documentation to include more errors, mitigations and descriptions here. We welcome additional feedback on any errors/metrics/logs etc that are not documented or available and what details would be useful here.
API Testing	What will the final value of epsilon be after the test period?	At this time, there is no plan to modify the epsilon value window. We encourage testers to experiment with different parameters and provide feedback.
Reporting	Report is getting generated, but also still getting PRIVACY_BUDGET_AUTHORIZATION_ERROR in return code.	We have provided guidance on specifying the correct reporting origin and aggregatable reports to avoid this error. We welcome additional feedback on the issue, in particular if there are recurring errors.
Key Discovery	What are the plans for the key discovery proposal?	We do not yet have a timeline for the roll out of the key discovery proposal but we are soliciting feedback from the ecosystem on the proposal here.

Customization	Seeking customization options available for the Aggregation Service.	Customizations of the Aggregation Service are not possible within the TEE but are possible for some components outside of the TEE. This is due to the privacy and security standards we need to maintain within the TEE. We are working on updating our documentation to help ad techs understand the architecture and what components are customizable. We would not be able to support certain customizations so we recommend ad techs to use our standard configurations where possible.
Spot vs. On-Demand Instances	Has the system been tested using spot instances versus on-demand instances? Are there any specific drawbacks to using spot instances, aside from their potential temporary unavailability?	We have not prioritized testing on spot instances because we recommend ad techs to use on-demand instances. The drawback of spot instances would be the job being interrupted during budget consumption. If the budget has been consumed and the job gets interrupted before the ad tech receives the summary report, ad techs would not simply be able to retry the job but would need to request budget recovery. Budget recovery is only recommended for catastrophic failures to preserve privacy. We recommend ad techs configure autoscaling to help minimize costs. Selecting 0 for autoscaling means there will be no running instances until a job is requested (note this may increase latency as instances will be spun up at the time of a request).
Known Errors & Solutions	Clarification needed regarding Aggregation Service job showing "Service Error"	This issue has been resolved here. We have also updated some of our error messages to make them more descriptive. For example, we have started throwing more specific permission/auth errors on AWS as opposed to previously when these errors were surfaced as internal errors. We have updated documentation on error codes and mitigation steps here and welcome additional feedback on errors that are not documented or available and what details would be useful here.

Private Aggregation API

Feedback Theme	Summary	Chrome Response
Key Design	Request for Private Aggregation key design guidance	While we do not have a Private Aggregation specific guide, both the Aggregation Service load testing framework and Advanced key management guide can be used as resources.
Contribution Budget	On what level is the contribution budget calculated and limited?	The contribution budget is 2^16 in a rolling 10 minute window and 2^20 in a rolling 24 hour window.

Limit Covert Tracking

User Agent Reduction/User Agent Client Hints

No feedback received this quarter.

IP Protection (formerly Gnatcatcher)

Feedback Theme	Summary	Chrome Response
Android	What is the plan to roll out IP Protection to Android?	While we are exploring bringing anti-covert tracking features to Android, including IP Protection, we don't have formal plans to share at this time.
API Usage	Question on if there is or will be an anti-fraud exception for IP Protection?	We strive to strike a balance between protecting users from being tracked across the web based on their IP addresses while minimizing disruption to the normal operations of servers, including the use of IP addresses for anti-abuse. While we cannot provide more details at the moment, we expect to provide them in the near future and look forward to continuing the discussion.
Bad Actor Identification	Concerns regarding the effectiveness of traditional security measures against malicious IP addresses.	IP Protection will not proxy 1P requests, so we expect most Intrusion Detection Systems will not be impacted. We plan to provide additional details in the future that address anti-fraud and site breakage concerns for incognito users.
IP Address Masking	If the news publisher site (1P) uses the same domain with the advertising site (3P), will the IP address be masked for both? If not, how does one distinguish the two?	IP Protection currently proposes a list-based approach to identify which third-party traffic goes through the proxies. However, even if an origin is on that list, it won't be proxied if accessed in a 1P context. We are finalizing the details regarding which specific 3P domains will

		be prioritized initially and how we'll precisely define 1P and 3P contexts for IP Protection.
IP Address Masking	Concern about IP protection and its impact on anti-fraud systems.	1Ps are not impacted by our IP Protection plans, so sites such as forums should have access to IP addresses for dispute resolution. We plan to provide additional details in the future that address advertising fraud concerns.
IP Address Masking	Is the IP masked in the header for impacted domains?	Users will be assigned to a geographic area based on their pre-proxy IP address representing their current location. You can find more details here.

Bounce Tracking Mitigation

Feedback Theme	**Summary**	**Chrome Response**
API Spec	Clarification needed on the behavior of Chrome's handling of extended navigation when a tab is closed.	We have resolved this issue here and updated the specification accordingly.
Nav Tracking	Discussion of a tracking mitigation approach involving a finite set of links to reduce entropy in cross-site requests.	We are continuing to discuss this topic with other browser vendors here, and welcome additional feedback and any specific proposals on this issue from the ecosystem.

Privacy Budget

As noted in the GitHub explainer and developer site, Privacy Budget is no longer being actively considered as part of the Privacy Sandbox proposals.

Strengthen cross-site privacy boundaries

Related Website Sets (formerly First-Party Sets)

Feedback Theme	**Summary**	**Chrome Response**
(Also reported in previous quarters) Related Website Set (RWS) Domain Limit	Request to increase the limit of Associated domains within RWS	Our response is similar to previous quarters: At present, we do not expect to increase the 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. If the user sessions are tied to login actions, we would recommend using the Federated Credential Management (FedCM) API to maintain functionality. We welcome feedback on other use cases which may be required here.
Cross-site cookie handling	Cross-site cookies set by a subdomain are not being passed in subsequent requests from the primary domain. The problem persists even with proper CORS and credential configurations.	We provided guidance here regarding correct usage of the requestStorageAccessFor API needing to specify the full origin (i.e. include subdomains) in order for cross-site cookies to be sent on subsequent fetch requests.
User Choice	Request for further information regarding requestStorageAccessFor used by RWS after the rollout of user choice, in particular how the implicit or explicit user gesture, which currently allows access to 3PCs, will function in the new system.	We expect that the behavior of RWS in either user choice mode, (i.e., regardless of whether users have chosen to retain or limit their 3PCs) will be consistent with existing/shipping behavior in Chrome with 3PCs allowed vs. 3PCs blocked with RWS enabled ("Allow related sites to see your activity in the group" setting). We recommend invoking hasStorageAccess first to check whether the embed already has access to unpartitioned cross-site cookies before invoking requestStorageAccess. The hasStorageAccess method will return true if the user chose to allow 3PCs. requestStorageAccessFor currently does not require a user gesture if 3PCs are allowed. We have opened a new GitHub issue to discuss and specify what the right behavior should be in this case, and welcome additional feedback from the ecosystem.

API Usage	Concerns about the lack of clarity regarding the use of RWS for "commercial" purposes, hindering their adoption. The stakeholder indicated interest in grouping publications for targeted advertising.	The intended use of RWS is to support core site functionality and core user journeys. We encourage using our purpose-built advertising APIs for use cases related to targeted advertising.
API Usage	Report of an issue with requestStorageAccess where they could set localStorage data but not cookies.	The issue was caused by a typo in the SameSite attribute. Ensure correct spelling and explicitly set it to None for 3PCs.
API Spec	Discrepancies in the JSON schemas across the repository, such as the misplacement of the “contact” field and suggestions for improvements, including the use of the “oneOf” keyword to ensure consistency.	We are investigating this feedback and will look into making these improvements to the schema in the near future. We welcome additional feedback here.
Third-party (3P) access to RWS	With given user consent, allow an outlet to indicate the 3P domains that will also have such access to the RWS API data.	Allowing 3Ps to combine their own unpartitioned data with RWS site data would undermine Privacy Sandbox's cross-site tracking protections. However, we are considering the potential for enabling 3Ps to maintain data partitioned to an RWS and are seeking feedback on the design for such a solution here.

Fenced Frames API

Feedback Theme	**Summary**	**Chrome Response**
API Question	Concerns on how Fenced Frames with no network access could break brand safety and fraud protection for advertisers.	We are tracking this concern in the context of our plan to enforce Fenced Frames. We will start looking soon into solutions that are compatible with Fenced Frames enforcement and as soon as proposals exist that are material enough we will share them.
API Question	Is Fenced Frames API still scheduled for 2026?	As stated in our public announcement, Fenced Frames will be required no sooner than 2026.

API Bug	When reportEvent() successfully sends click data from a cross-origin subframe, setReportEventDataForAutomaticBeacons() does not overwrite the top frame's data.	We are discussing this issue and welcome additional feedback here.

Shared Storage API

Feedback Theme	Summary	Chrome Response
App Advertising	There is no equivalent of the Shared Storage API in Privacy Sandbox on Android.	We are evaluating solutions on Android based on use case needs and platform constraints. We do not have any plans to share at the moment, but we welcome additional feedback from the ecosystem on this issue.
API Usage	Confusion regarding the need for an additional javascript worklet for implementation for Shared Storage API.	We are investigating this feedback and looking into potentially updating our documentation to explain the need for additional worklet scripts for Share Storage API.
Unreliability	Shared Storage API could change significantly or be dropped based on the privacy criticisms, making it an unreliable base to build on.	We do not have plans to significantly change or drop the Shared Storage infrastructure. The main changes that have been announced are to the Select URL output gate where fenced frames will be required and event level reporting will be deprecated. However these changes will not go into effect until at least 2026.

CHIPS

Feedback Theme	Summary	Chrome Response
Partitioned Cookies	Chrome does not differentiate partition keys based on frame ancestors, unlike Firefox, leading to inconsistencies.	Chrome adopted the "ancestor chain bit" to resolve the security concern and discrepancy with Firefox's behavior. We have set this out in further detail here.
Partitioned Cookies	Embedded iframes with different storage access levels share and overwrite	For this particular configuration, our recommendation is to use unpartitioned cookies with an invocation of Storage Access API.

	the same partitioned cookie, causing inconsistencies in authentication states.	We have discussed this in further detail here.
Partitioned Cookies	Will partitioned cookie jars be cleared when 3PCs are disabled? Is there a way to test this behavior?	We do not have any plans to share at this time. However, developers can test this functionality by manually deleting partitioned cookies via the Chrome DevTools Application>Cookies panel.

FedCM

Feedback Theme	**Summary**	**Chrome Response**
Identity Provider (IdP) Registration Scope & Organization Chooser	Question on extending IdP registration from the current same-origin policy to a same-site policy. This change would allow broader and more flexible identity management, such as enabling a university's welcome page to register multiple subdomain-based identity providers without needing separate origin-specific registrations. Feedback on improving debuggability, handling approved clients for silent mediation, clarifying cookie behavior, allowing customization of the popup wording, and addressing timeout issues.	We acknowledge this feedback and are considering how an organization chooser could be incorporated into FedCM. We welcome additional feedback from the ecosystem here as we continue to refine this approach.
IdP Cookies	Discussion on the impact of implementing short-lived session cookies as part of the Device Bound Session Credentials (DBSC) proposal. Concerns are raised about user experience in FedCM, where expired IdP cookies require a user-visible modal for renewal.	The proposed DBSC should allow for cookie renewal without user interaction, ensuring continuous user experience. We have discussed this issue in further detail here.
API Spec	Question on appropriateness of using "NetworkError" in	We agree that "TypeError" would be more appropriate for this situation since it reflects a

	the FedCM API specification when the size of the "providers" array is not equal to 1.	coding error rather than a network issue. We will consider this change and explore the possibility of removing this restriction in future updates as we progress towards multi-IdP support. We welcome additional feedback here.

Fight spam and fraud

Private State Token API (and other APIs)

Feedback Theme	**Summary**	**Chrome Response**
Deprecation Trial & Mode B	Concerns about the deprecation trial, Mode B testing, the potential discontinuation of Private State Tokens (PSTs), and the possibility of a new API better suited for their anti-fraud use case.	The deprecation trial and Mode B testing remain unchanged. We will communicate any updates through the dev blog. We have no plans to discontinue PSTs and we are discussing ongoing feature development and updates on GitHub. We have not announced any new APIs, but we welcome feedback on how PSTs can better address anti-fraud use cases.
API Feedback	Request for the capability of revoking tokens to address an anti-fraud use case.	While the issuer could revoke all tokens by changing the keys they use, individual token revocation is infeasible with the API as it would require allowing the issuer to associate token issuance and redemption which breaks the PST privacy model of preventing tracking between the two events.

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 July 22, 2024, Google's display ads platforms on both the buyside and sellside published results from their experiment testing privacy-preserving solutions and Chrome's Privacy Sandbox APIs in combination (Topics, PA API and ARA) via Chrome-facilitated testing. The buyside results can be seen here, and the sellside results can be seen here: Ad Manager and AdSense. The full version of the whitepaper on Chrome-facilitated testing is 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. Google is working with the CMA to resolve these concerns, following the process set out in paragraph 17(a)(ii) of the Commitments. The concerns are summarized in the CMA's quarterly update report. 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 as set out below.

Stakeholder concerns

The CMA has shared with Google certain concerns expressed by stakeholders. The concerns set out below are not exhaustive, and are in addition to those addressed above.

Competition Feedback – The CMA has shared stakeholder feedback relating to Google's market power and a stakeholder concern that the Privacy Sandbox proposals could be anti-competitive or advantage Google, in particular through the use of 1P data. The CMA has also shared a stakeholder concern that the Privacy Sandbox may discriminate against 3Ps by impacting signals used by 3Ps for validating performance. As was set out in Google's Q1 2024 progress report, Google has committed to design and implement the Privacy Sandbox proposals in a way that does not distort competition by self-preferencing Google's own business, and to take into account impact on competition in digital advertising and on publishers and advertisers, regardless of their size. We continue to work closely with the CMA to ensure that our work complies with the Commitments, and we welcome feedback on how the APIs perform for different types of stakeholders.

The CMA has also shared feedback from a stakeholder that considers GAM's decision to not participate in PA API component auctions unless it is a top-level seller to be a form of technological tying. Our response is unchanged from previous quarters:

"Response provided by Google Ad Manager:
Google Ad Manager's plans for the Protected Audience API do not include supporting Google's publisher ad server without the control of the top-level Protected Audience 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 Protected Audience 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 functionalities 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. In short, therefore, we do not view a publisher ad server's activity of running the top-level Protected Audience auction as distinct from the other activities of the publisher ad server."

Google has received stakeholder feedback from the CMA that the only cross-domain and cross-device identification solutions available will be those based on email addresses. The stakeholder considers that this is anti-competitive and unfair. This stakeholder feedback is based on a misunderstanding of the facts. There are numerous signals and technologies that enable effective targeting and measurement of online advertising, including solutions that facilitate cross-site tracking independent of 3PCs, not all of which are based on email addresses. 3Ps have developed solutions based on signals such as publisher-provided information and contextual information. Moreover, IP addresses are set to remain available for users in default browsing mode, under Google's announced plans to introduce IP Protection in Chrome's Incognito mode. Ultimately, there is ample opportunity for developers to build privacy-enhancing technology solutions for cross-domain and cross-device targeting and measurement on top of the building blocks we're offering as well as non-Privacy Sandbox building blocks.

Ads relevance - The CMA has shared a stakeholder concern that the Privacy Sandbox could reduce ads personalisation in open display advertising, thereby reducing the relevance of ads displayed to users. The Privacy Sandbox APIs are not intended to be direct, one-to-one replacements for all 3PC-based use cases or to be a standalone ad tech solution. Instead, they are designed to provide foundational elements that support core business objectives for marketers and publishers (like driving online sales and serving relevant ads), without relying on cross-site identifiers. Developers can utilize them alongside other technologies and inputs to achieve those outcomes. Google has previously shared a blog post on Maximizing ad relevance to educate the ecosystem about maximizing performance using a range of privacy-safe signals.

We believe the current Privacy Sandbox APIs – generally available in Chrome since September 2023 – are ready to carry the ecosystem into a more private future. And we're committed to pushing privacy-preserving technologies forward for years to come, both in terms of privacy and utility.

Privacy Feedback – The CMA has shared with Google stakeholder feedback that Google's quarterly reports to the CMA do not discuss how the Privacy Sandbox APIs improve privacy for individuals. Google's quarterly reports provide the CMA and ecosystem with an update on Google's progress on the Privacy Sandbox proposals; updated timing expectations; substantive explanations of how Google has taken into account observations made by 3Ps; and a summary of interactions between Google and the CMA.

We have set out in detail how the Privacy Sandbox APIs provide a higher level of privacy as compared to 3PCs, enabling key advertising use cases without sharing the user's identity with 3Ps, in our blog post on 'How Privacy Sandbox raises the bar for ads privacy'. We also frequently respond to stakeholder feedback relating to user privacy in our quarterly reports, and we welcome feedback from the ecosystem on any aspect of the Privacy Sandbox, including those relating to user privacy, here.

Attribution Reporting API – The CMA has shared feedback that the re-identification risk relating to the Attribution Reporting API can be addressed through more proportionate measures such as reliance on de-identified data storage. The stakeholder considers that there is no reason to ban all event-level data points. The UX needs to be made transparent so that the resulting data trades are transparent and fair. The stakeholder considers that the re-identification risk is not unique to 3Ps, and that it also arises within Google.

We have carefully considered the proportionality of the privacy protections included in our API design, and we have no current plans to ban event-level data points in the Attribution Reporting API. The design of our UX is focused on providing adequate transparency and controls for Chrome users, by clearly communicating information about the functionality and purpose of the Attribution Reporting API, and empowering users with controls over its availability.

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.

[redacted]

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 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 Signed by: [redacted]..

Full name [redacted]..

Date [redacted]..

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