Privacy Sandbox Progress Report

Q1 Reporting Period – January to March 2023

Prepared for the CMA, 21 April 2023

Overview

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

Progress of Privacy Sandbox Proposals

Google has been keeping the CMA updated on progress with the Privacy Sandbox proposals in its regular Status Meetings scheduled in accordance with paragraph 17(b) of the Commitments. Additionally, the team maintains the overall Privacy Sandbox developer documentation with specific pages for each API, an overall status page, along with regular updates for the Relevance and measurement unified origin trial. Key updates are shared under the "Privacy" tag on the developer blog along with targeted updates shared to the individual developer mailing lists.

Updated Timing Expectations

Google's latest expectations for the timing of the Privacy Sandbox proposals are set out in the Privacy Sandbox Timeline.[1] The summary below includes all Q1 2023 updates, covering the period from January 1 to March 31, 2023. Google is working towards the removal of third-party cookies in H2 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 Q1 2023 Timeline Updates	
January Timeline Updates	• Updated timing of CHIPS rollout on Chrome stable from “January 2023” to “February 2023”
February Timeline Updates	• No updates
March Timeline Updates	• No updates

After consultation with the CMA, Google’s Privacy Sandbox Team is planning to publish updated guidance on testing the Privacy Sandbox proposals, including the suggested timeframe for testing, in the course of Q2 2023. For more details about “Google Roadmap for Effectiveness Testing of the Privacy Sandbox Proposals”, see relevant section below.

Taking into account observations made by third parties

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

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

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

The explanations of Chrome’s responses to feedback were developed from published FAQs, actual responses made to issues raised by stakeholders, and determining a position specifically for the purposes of this public reporting exercise. Reflecting the current focus of development and testing, questions and feedback were received in particular with respect to Topics, FLEDGE and Attribution Reporting APIs and technologies.

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

Glossary of acronyms.

CHIPS - Cookies Having Independent Partitioned State
DSP - Demand-side Platform
FedCM - Federated Credential Management
FPS - First-Party Sets
IAB - Interactive Advertising Bureau
IDP - Identity Provider
IETF - Internet Engineering Task Force
IP - Internet Protocol address
openRTB - Real-time bidding
OT - Origin Trial
PatCG - Private Advertising Technology Community Group
RP - Relying Party
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**
Testing and trialling	Relevance of the testing to inform the CMA's assessment if the Privacy Sandbox APIs are not completed by the time the test begins	The development of the Privacy Sandbox APIs is progressing at pace. They are already available in OT for testing and will be generally available for 100% of traffic this summer. Additionally, we have clarified timelines for certain features (e.g. FLEDGE event-level reporting, FLEDGE rendering with iFrames) that will not be impacted sooner than 2026. We encourage the ecosystem to test the APIs and provide feedback to the CMA based on what testers expect to rely on once third-party cookies are deprecated. This can contribute to their assessment of the likely impact of third-party cookie deprecation.
User Controls	Clear guidance to ecosystem on user controls implications of Sandbox APIs	We can't provide legal advice on what user controls the ecosystem can use. At the same time, Chrome is experimenting with showing updated Privacy Sandbox ("Enhanced Ad

		Privacy") user controls to a very small percentage of users, as part of our ongoing effort to improve the Privacy Sandbox technologies. The updates include clearer, more helpful language and layouts. Once Chrome has evaluated these refinements and decided whether to expand them to a larger population, we can share more information with the ecosystem.
Data leakage	Risk of first-party data leakage to Google and other parties in the event the browser is compromised	Our FLEDGE explainer makes it evident that one ad tech's data is only shared with that same ad tech (either with their worklets or their trusted servers) or when explicitly shared by that ad tech (e.g. a buyer shows a seller the ad URL they want to display). The one exception being that the k-anonymity check must be done by a global centralized server, which is an area we continue to devote significant resources to. Please see the k-anonymity explainer for a detailed view of how we're thinking about privacy. Beyond that, we are open to providing more details on how the ad-tech protections employed in the design of the k-anonymity server work.
Additional forum for discussion	Request for additional forum to the W3C for non-technical ecosystem players to share feedback.	The Privacy Sandbox feedback form is appropriate for general and specific comments, technical and non-technical. The Improving Web Advertising Business Group is a forum for discussion via weekly calls and a GitHub repo. The Privacy Sandbox Feedback page on developer.chrome.com explains other mechanisms for providing feedback and engaging in discussion. Chrome also continues to hold events like public Office Hours to facilitate question asking and content sharing. In addition, Chrome has hosted or attended more than a dozen industry events this past quarter.
Timeline clarification	Clarification on the exact date for General Availability in Q3 2023	Per the timeline published on PrivacySandbox.com, General Availability is targeting to begin rollout with the release of Chrome version 115.
reCAPTCHA	Impact of Sandbox APIs for reCATPCHA's spam detection use	We get feedback from reCAPTCHA periodically to ensure Privacy Sandbox

	case	proposals do not significantly impact web safety or fraud. They are developing their own plan to prepare and adjust for third-party cookie deprecation, and so this question is best fielded by them.
Chrome Extensions	Will Privacy Sandbox technologies like Anti Covert Tracking (ACT) measures apply to Chrome extensions?	We have not made any announcements on whether ACT may apply to Chrome extensions. However, if a technology covertly gathers information on a user, this would not align with our privacy principle.

Show Relevant Content & Ads

Topics

Feedback Theme	**Summary**	**Chrome Response**
TAG Design Review	TAG released Early Design Review of Topics.	We remain committed to Topics and have shared an update on our commitment to Topics here and here. We responded, point-by-point, to the TAG review here and shared our higher level vision here. The Topics API will remain part of the collection of APIs that the ads ecosystem should test during 2023 — and we hope the testing feedback we hear, and the implementer experience we gain, will be valuable contributions in future work towards cross-browser standards work in this space. We look forward to continuing to engage with the ecosystem on how to ease a possible transition where Topics API could be an agreed upon standard with cross-browser compatibility.
Approach to Topics	Support for the open approach Chrome has to developing the Topics API.	We appreciate the sentiment and we look forward to continuing working with the industry group to develop a Topics API that provides value to the ecosystem overall.
(Also reported in Q3 2022) Topics taxonomy not granular enough	Broad Topics taxonomy does not include more granular topics, including region specific.	Improvements to the taxonomy are an ongoing effort, and in Q2 we will announce an updated taxonomy for the Topics API. To craft this new taxonomy, we worked closely with companies from across the ecosystem. We are actively seeking feedback on the

		taxonomy that would be most useful for the ecosystem. In evaluating whether to expand the number of topics or include more granular topics, there are a few considerations including 1) potential privacy implications (e.g. more topics may introduce fingerprinting risk) and 2) ability to retrieve previously observed topics (e.g. with more topics, there may be less of a chance that an ad-tech has seen the chosen topic in the past).
(Also reported in Q4 2022) Impact on first-party signals	Topics signal may be highly valuable and as a result devalues other first-party interest-based signals.	We believe interest-based advertising is an important use case for the web, and Topics is designed to support that use case. We understand that some large publishers are concerned that Topics will negatively impact their first-party data strategies. We look forward to ecosystem testing which will provide insights into the impact Topics has on publishers.
Non-Ads related Topics use cases	Use of Topics for purposes other than displaying interest-based advertising.	Topics is designed to address the interest-based advertising use case, which we believe is a critical use case for the free and open web. We are currently seeking feedback on other use cases and are evaluating them.
Default Opt In Status	Regional legislation impacts for Topics consent default	It is not our position to comment on legal opinions.
(Also reported in Q3 2022) Miscategorized sites	Ads targeting when topics miscategorized for a given site	In Q2 we will announce an updated classifier for the Topics API and look forward to engaging with the ecosystem on it. In response to the current feedback, sites are classified through a combination of a human-curated override list, containing the most popular sites, and an on-device ML 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-labelling as a method to encode different (and potentially sensitive) meanings into topics; sites misrepresenting their topics for financial gain; 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.
Topics Classifier	Request to return additional information showing the reasons when "No Topics" is returned to the caller for debugging purposes.	We understand and appreciate that debugging tools are helpful to developers, as they work to integrate Topics API into their systems. However, by exposing additional information (such as the reason why no Topics were returned), we may inadvertently share information that enables parties to uncover additional details (e.g., if a user is in incognito mode, has disabled the API, etc.) beyond what is intended, harming user privacy. While we don't plan to provide additional debugging tools at this time, we are open to feedback about which tools would be valuable.
Private Information Retrieval (PIR)	Request for Topics API to adopt PIR.	We have previously investigated using PIR and have shared the trade offs here.
Bid Stream	Will Topics be represented distinctly from Seller-Defined Audiences in the bid stream?	The Topics API is a Privacy Sandbox proposal developed by Chrome, which is distinct from the IAB Tech Lab's Seller-Defined Audiences proposal. We expect the two to be represented distinctly within the bid stream. You can see how topics will be represented in OpenRTB bid requests here.

Protected Audience API (formerly FLEDGE)

Feedback Theme	**Summary**	**Chrome Response**
FLEDGE feature availability	Clarification on timelines for testing and implementation for FLEDGE features such as Fence Frame enforcement, K-Anonymity etc.	We have shared a blog post of various scoped FLEDGE features and when they will be supported. We welcome additional feedback on the announcement as we continue to develop FLEDGE.
Product rendering restrictions	Request to loosen Ads Composed of Multiple Pieces restrictions for FLEDGE Fenced Frames	As we announced in February, usage of Fenced Frames will remain optional until at least 2026, and iFrames behavior will be

		supported by urn-iframes. We welcome further discussion on this topic.
Scalability Issues	FLEDGE performance as usage scales	We are actively following up on the feedback and understanding more context so that we may propose actionable solutions. The first step was to separate the feedback into two categories, which we have done: 1. SSP-driven filtering to optimize queries-per-second (QPS) load on both a) themselves and b) the DSPs 2. Interest group DailyUpdate logic to optimize QPS load on DSPs
(Also reported in Q3 2022) Visibility of bidding logic	Concern that DSP bidding logic will be exposed in JavaScript.	Q1 Update: We have shared a proposal here that would limit the ability of adversaries to request data from the server in an exploratory (force browsing) fashion and we welcome ecosystem players to share their feedback or support for the proposal.
Testing difficulties	Ability for smaller DSPs to properly test FLEDGE, and risk advertisers are only interested in testing with larger DSPs.	We are committed to working with smaller DSPs and strongly encourage expanded testing among DSPs and advertisers of all sizes as FLEDGE moves to general availability. We would be interested in hearing how we can best assist them in testing FLEDGE with others in the ecosystem, and welcome ideas and industry efforts to motivate advertisers to test with smaller DSPs.
Dynamic remarketing	Will Dynamic remarketing still be possible with FLEDGE post-third-party cookie deprecation?	We are considering a response to this question and welcome if ecosystem players can share additional insights on how they intend to use Dynamic remarketing.
Fraud/Abuse	How can the ecosystem reduce the risks/stop bad actors or buyers from positioning themselves as a desirable audience?	We look forward to engaging with ecosystem players further on fraud and abuse, and welcome more feedback on this area.
User preferences	Process to save user preferences and use in ad selection.	For specific ads, the relevant ad tech is the party best positioned to offer controls over which creatives are shown or how they are selected.
Quantitative Testing Proposal	For Quantitative Testing to be fair, should the test be conducted on traffic without third-party cookies or with SSPs that only use FLEDGE?	We appreciate this feedback and are working together with the CMA to design experiments that will provide a reliable picture of the impact of third-party cookie

	How can the mixing of signals from third-party cookies be avoided?	deprecation and the introduction of the Privacy Sandbox proposals on the ecosystem. We encourage additional feedback on the CMA's Quantitative Testing proposal to be shared directly with the CMA.
Clearer documentation	Request for clearer documentation on auction configuration.	We are hoping to share a blog post with additional overview on FLEDGE Auction Reporting in the coming weeks.
Parallelization	Will the Bidding and Auction (B&A) Service support Parallelization?	An ad tech using B&A servers can start multiple servers that can serve results in parallel.
Abuse Mitigation	Will the FLEDGE K-anonymity server using Private State Tokens be enough to ensure user privacy?	The motivation for k-anonymity is less focused on microtargeting and more on having some backstop during the interim phase where FLEDGE allows event-level reporting. We have shared more thoughts here and welcome additional feedback.
ES Module conflict	Request to drop generateBid as a global function as it restricts with the ES module.	We are discussing this request and welcome additional feedback.
Component Auction	Request for Publishers to have more control over auction designs	B&A plan to support component auction, same as Chrome on-device.
B&A Timelines	Clarity on the timeline for ad techs interested in testing B&A Servers	We just updated the B&A Explainer and we updated the text in the Timeline section to include clear definitions of timelines for different phases of Chrome-B&A testing, after aligning with the CMA.
Time out control scheme	Enhancing the time out control scheme currently available for FLEDGE	This is an interesting proposal. We will add this to the queue of proposals to study, and report on our developments.
Creative Bidstreams	Ability to review, and filter a winning bid, based on the creative.	This is an interesting proposal. We will add this to the queue of proposals to study, and report on our developments.
reportWin	Proposal to provide additional information on highest scoring bid from a different interest group owner other than the winner in the reportWin function.	This is an interesting proposal. We will consider adding additional signals in aggregate reporting and welcome additional feedback here.
Event Types	Standardizing event types across measurement APIs when integrated with FLEDGE	This is an interesting proposal. We will add this to the queue of proposals to study, and report on our developments. It will require coordination with our broader efforts in this field, as it would affect other Privacy Sandbox APIs beyond FLEDGE. We welcome additional feedback here.

Long term solutions for Event Level Reporting	Interest in keeping certain data such as highestScoringOtherBid available even after third-party cookie deprecation.	As we shared in the February blog post, Event-level auction win reporting will be supported until "at least 2026". We do not have further details to share at the moment but we welcome additional feedback on why it is important to keep certain data available after third-party cookie deprecation.
Interest Groups Limit	What is the limit to the number of interest groups that an origin can add a single browser to?	Chrome allows up to 1000 interest groups per owner, and up to 1000 interest group owners. These are meant to be guard rails, not to be hit in regular operation.
Event-Level Signals	Support for a proposal to have event-level signals for generateBid and reportWin which could be used in machine learning training.	We have shared our decision for browser designed signals and ad tech defined signals here and welcome additional feedback.
Bidding Script	Include user ID in the url to the bidding script.	This will not be possible as FLEDGE has the additional requirement that the tuple of the interest group owner, bidding script URL, and rendered creative must be k-anonymous for an ad to be shown.
K-anonymity enforcement	Is k-anonymity enforced on (componentAd, size) pair?	Yes, it will be. See here.
B&A Services Requirements	How do B&A Services support participants integrating with the on-device FLEDGE and others with B&A services?	We are still finalizing the design and welcome additional feedback here.
Post-view attribution	Will Post-view attribution be supported	Currently we don't have any kind of standard definition of viewability and rely on the creative itself to mark a view event. See here.
Lookalike targeting	Can Privacy Sandbox support "lookalike targeting"?	We are discussing the use case here and welcome additional input.
Real time monitoring API	Proposal for a Real Time FLEDGE monitoring approach.	We are discussing the proposal and welcome additional input here.
FLEDGE reporting	reportWin and reportResult should be made in random order to prevent over or under reporting.	reportResult() needs to be executed first by the seller before reportWin() by the buyer so that seller signal from reportResult() can be included in reportWin(). Please see the explainer for more information.
Custom Key Value (K/V) Servers	Will custom K/V Servers be supported in the future?	We are discussing the question here and welcome any additional input.
Top-Level Auction	Does one have to be the ad server to run top-level auction mechanics?	The FLEDGE API does not specify which party must call it – there are no requirements in that sense in the design of

		FLEDGE. Anyone can run the FLEDGE auction (including multi-seller auctions). As mentioned in the Q4 2022 report, FLEDGE allows each publisher to choose the structure of the auction, including the choice of top-level and component sellers.
API Scope	Does FLEDGE intend to work with first-party data?	We will publish content in Q2 2023 clarifying that first-party data is indeed usable with FLEDGE to both 1) use as logic to determine interest group membership, and 2) to feed as user bidding signals for use in subsequent bidding logic generation.
Cross domain interest groups	Possibility of creating cross-domain interest groups	Any information available at the time of adding a browser to an interest group can be used to inform that audience. When third-party cookies are phased out, the availability of cross-site data to inform interest group creation will be limited.
Client side bidding logic	How to port existing server side bidding logic to Client Side?	We are interested in learning more on what areas are challenging or currently lacking in the porting process, and welcome any additional feedback or insights here.
K/V Server Value	Do K/V Server values need to be in String?	The value needs to be in string but they can store objects in JSON or protocol buffer and serialize them into string.
Advertiser blocklist	Which signals would be appropriate to provide a buyer for an advertiser blocklist?	Appropriate place is either in auctionSignals or in perBuyerSignals.
Bidding unit	Support for different bidding units such as CPI and CPM	We would like to learn more about why this is needed given the current design and would welcome additional feedback.
Auction Logic	Does the browser or the Ad server decide the winner of an auction?	All winner selection is executed inside the sandbox, and all decisions are made by the seller's code. The browser simply provides a sealed, private environment, inside which buyer and seller code runs.
Permissions-Policy	Will the current FLEDGE Permissions-Policy continue to be enforced after the OT has ended?	For the OT, the current default allow lists of both features are temporary and will be changed. We are interested in hearing how long ad techs will need to prepare for the change before we begin to enforce it.
Signal size constraint	Trusted Bidding Signals requests are coalesced across multiple interest groups with the same trustedBiddingSignalsUrl, the 2MB size limit is a constraint.	The constraint exists for on-device callers to prevent overwhelming resources on the device. Callers from a B&A Server will have a more relaxed constraint, as discussed here.

Reporting signals	Add an additional signal script-errors to allow for the retrieval of the number of client side errors per interest group owner and per computeBid or reportWin/reportResult.	We are considering potential privacy concerns of this proposal and welcome additional insights from ecosystem players on why this is needed.
K-Anonymity Window Size	Increase the K-Anonymity Window Size from the current 7 days limit.	This is under consideration and we are currently awaiting and welcome additional input from the ecosystem.
Device Performance	How does FLEDGE handle device performance if the user is in a large number of interest groups?	FLEDGE offers several timeout, prioritization and limit options across SSPs and DSPs that give ad tech fine grained control in situations where device performance may be one reason to limit auction participation in a situation where the device is in a large amount of interest groups.
B&A Services testing	Request for ecosystem players to use their own server during the testing phase in order to have more logs available for debugging.	B&A allows users to launch and scale the servers from approved cloud providers. To maintain user privacy, we enforce execution to be done within a TEE. We are going to release an explainer about debugging of B&A TEE soon and are developing features to support that. We are seeking additional feedback on the topic.
Regulatory Requirements	Will FLEDGE work with cloud providers in different countries to support compliance with local regulatory requirements?	We are always open to suggestions for other cloud providers, but currently we are planning at least to support GCP and AWS when third-party cookie deprecation is enforced. See this explainer for more information.

Measuring Digital Ads

Attribution Reporting (and other APIs)

Feedback Theme	**Summary**	**Chrome Response**
Noise impact data analysis	Guidance on how to perform data analysis on the impact of Noise.	We have shared additional documentation regarding noise and design decisions that can be used to change the impact of noise on ad tech data. A more detailed guide is available here.

Null reporting	Clarity on the implementation of null reports.	We are currently working on a proposal for implementing null reports and will have more details to share soon. Implementing null reports will allow us to reduce report delays without compromising privacy.
Noise Level	Adjusting the noise level based on attribution window length.	We welcome this proposal and are looking into adding it to the specification, and welcome additional feedback about it here.
Trigger Data Size	Why is the trigger data size limited to 3 bits?	The size is limited to 3 bits and 8 distinct values to ensure that the amount of cross-site/contextual information about a user is limited. We welcome ecosystem players to submit feedback on whether the current parametrization for event-level reporting makes sense here.
Event-Level Reporting Triggers	Allowing prioritization within a deduplication key.	We are exploring solutions to this problem here and welcome additional input.
Debugging Support	Clarity on debugging post-third-party cookie deprecation.	We would like to support debugging after third-party cookie deprecation and are considering possible options. We are seeking additional feedback and ideas here.
Click Through Conversion Alternatives	Request for more guidance on alternatives for click through conversions.	We encourage the ecosystem to use Attribution Reporting API as a durable private measurement system for applicable conversion measurement use-cases. Other alternatives exist and ad tech providers will need to decide the appropriate solution based on their desired privacy and utility needs.
Billing use cases	Clarity on the extent Attribution Reporting will support conversion-based Billing use cases.	We are working on posting publicly clarifying the scope of the Attribution Reporting API for billing. Attribution Reporting API was not initially scoped in a way that directly supports CPA billing, it supports CPC and CPM billing which is the billing structure the majority of ad techs use. This is something we may support in the future if there is additional ecosystem feedback
Use Case Support	Use case documentations for measurement API.	We are working on clarifying the documentation for all Privacy Sandbox reporting surfaces.
Click quality	Request to add signal to distinguish	We are discussing the request here and

	intentional and unintentional clicks on an ad.	welcome additional input.
Measurement solution	Support for measurement solutions across multiple DSPs.	The Attribution Reporting API can be used by measurement providers to dedupe between multiple DSPs. Additionally, we are proposing supporting a list of urls in attributionsrc which will make it easier for DSPs to support measurement provider Attribution Reporting API requests. We welcome any additional feedback on the proposal above.
Event-Level Reporting	Request to have the number of days before the report is sent directly available.	This request can already be calculated by ad techs using the current information available. We have not heard any other ecosystem feedback regarding this request, but we are open to feedback on it.
source_registration_time	Add source_registration_time in Event-level Attribution Reporting API.	We are considering this request and welcome additional feedback on whether the ecosystem players find it a useful feature to have.
Incognito Mode	Will measurement solutions be available when the user is in Incognito Mode.	No, measurement solutions will not be available when a user is in Incognito Mode. Third-party cookies are off by default in Incognito Mode.
Data Clean Rooms	Will Measurement APIs be compatible with clean rooms?	A typical data clean room is an environment where individual identifier data from different sources are uploaded into a database to run analyses based on merging that underlying data. The two measurement frameworks for Privacy Sandbox APIs are event-level reports and summary reports. Event-level reports do contain an ad-tech provided event-ID that could be used in a data clean room, but the conversion side information associated will be limited and noisy. Encrypted aggregatable reports cannot be used directly in a clean room, but the summary results provided by the aggregation service could be used as an input to analyses you perform or as supplemental information.

Aggregation Service

Feedback Theme	Summary	Chrome Response
(Also reported in Q4 2022) Reporting delays	What is the expected reporting delay?	Q1 2023 Update: Following partner feedback, we have shared proposals to decrease delay and to mitigate the impact of delay. Both proposals have been supported by ad techs during WICG calls.
No duplicates rule	How to handle a “delayed aggregatable report” if aggregatable reports, which have the same shared ID, were already processed?	We have shared a proposal on adding ‘extra report delay’ to an aggregatable report’s shared information and the definition of shared ID for Aggregate API to partially address the impact of delay loss on Aggregate API. We welcome any feedback on the proposal.
Data processing	Request to enable support for multiple passes of data while respecting differential privacy, using Privacy Budget.	We are discussing the possibility of using a more flexible way of consuming Privacy Budget to enable this use case and welcome additional feedback here.
(Also reported in Q2 2022) Query Ergonomics	Enable querying aggregate of keys.	Q1 2023 Update: The feature request is still being considered but we do not have any proposals to share at this time.
OT Limitations	Clarify the scope of Aggregation Service such as the "no duplicates rule" which is not currently applied in OT.	We are looking into updating our documentation to clarify what will be available in OT vs GA.

Private Aggregation API

Feedback Theme	Summary	Chrome Response
Private Aggregation Contribution Budget	The L1 contribution budget is too restrictive.	Each call to the Private Aggregation API is called a contribution. To protect user privacy, the number of contributions which can be collected from an individual are limited. When you sum all aggregatable values across all aggregation keys, the sum must be less than the contribution budget.

		Under the current design, we set a limit on the contributions for a particular reporting origin over the last ~24 hours (as a rolling window). That's the L1 contribution budget mentioned in the feedback. We do suggest that developers scale the values they contribute based on expected volume (i.e. not just using a value of 1). So, it might make sense to use a smaller value for the more common events to avoid exhausting the budget. We're currently seeking some feedback on the Private Aggregation API's contribution budget on both the numeric bound and the scope. We are considering moving the scope from per-origin to per-site and moving the existing bound to a ten minute window with a larger daily bound.

Limit Covert Tracking

User-Agent Reduction/User-Agent Client-Hints

Feedback Theme	**Summary**	**Chrome Response**
UAR Adoption	Of the top 10,000 sites in the UK, only 1% of sites using programmatic advertising are sending HTTP client hints. DSPs who have not migrated may see an impact on anti-fraud capabilities.	After running an analysis on the same data set, we have found that if you account for UA-CH usage via HTML <meta> tag, and the JavaScript APIs, the number of sites using UA-CH is significantly higher than the 1% figure provided in the feedback. Based on this, and other facts including ecosystem feedback, we feel confident moving forward with the gradual rollout of Phase 6 of UA Reduction, in accordance with the published timeline, while keeping the CMA informed. We note that sites have had nearly two years of lead time to prepare for the transition and a deprecation trial is still available for sites that feel they are not ready.

Hints for additional form factors	Request for UA-CH to provide additional form factors such as TV, VR.	We welcome this proposal and are looking into incorporating it to the design. We welcome additional feedback here.
Automated Testing	Request to resolve UA-CH bug in headless Chrome before UAR Phase 6 is shipped.	The bug in question has been fixed.
UA-CH support on iOS	A site that relies on granular UA info for ads use cases notes that Chrome on iOS is not supported.	For non-Safari iOS browsers (including Chrome on iOS), the WebKit project will need to add support for UA-CH before they can be enabled (because they control the network stack).

IP Protection (formerly Gnatcatcher)

Feedback Theme	**Summary**	**Chrome Response**
(Also reported in Q4) Geolocation use cases	IP Protection may prevent legitimate geolocation use cases from working in the future, such as content personalisation based on geolocation.	Our response is unchanged from Q4 2022: *"We are working with stakeholders to ensure that Chrome continues to support legitimate use-cases for IP addresses. We are seeking ecosystem feedback on IP Geolocation granularity here."*
Regulatory Compliance	If a region has under 1M population, the current threshold of 1M for IP Protection would prevent websites from using IP addresses for regulatory compliance.	We are working with stakeholders to ensure that Chrome continues to support legitimate use-cases for IP addresses. We are seeking ecosystem feedback on regulatory compliance on IP Protection.
Abuse Mitigation	Parties can circumvent IP Protection by sharing unmasked IP addresses to others.	We are conscious of the risk that the current IP Protection proposal might not technically prevent parties from sharing unmasked IP addresses with others. We are working on mitigations that will avoid this risk of abuse. As we iterate on the proposal, we encourage more feedback and discussion. Specifically, we would like to know of any use cases where parties believe they need to share unmasked IP addresses with other parties.
Network Blocking	Parties can circumvent network blocking by using IP Protection Proxies.	The entity performing the block will need to disable IP Protection for this scenario. We have responded to the issue here and welcome additional feedback.

IP Address Block Lists Impacted by IP Protection Proposal	Many ad tech companies utilize a basic block list of IP addresses, such as the TAG Data Center IP list, to prevent bidding on ad inventory that is highly likely to be fraudulent (or at least not monetizable). In the event an ad tech is also a tracker and could be subject to the IP Protection proposal, that company may lose the ability to perform a basic check against ads prior to purchasing advertising inventory.	We encourage more feedback and discussion on the IP Protection Proposal on potential issues and solutions. One option is applying similar such lists to IP Protection, such that we are not proxying clients originating from previously flagged IP addresses.

Strengthen cross-site privacy boundaries

First-Party Sets

Feedback Theme	**Summary**	**Chrome Response**
(Also reported in Q4) Domain limit	Request to expand the number of associated domains.	Our response is unchanged from Q4 2022: "*We have clarified in WICG calls that Chrome is committed to providing a usable solution that considers users' privacy interests as well. In that vein, we would appreciate feedback from the community on specific use cases that may be impacted by the domain limit, so that the team can consider ways to address these use cases while continuing to protect user privacy.*"
Alternative FPS submission	Proposal for alternative way to submit global lists for FPS.	At this time, we are preparing to ship FPS in Chrome, and have set up a centralized GitHub repository to accept set submissions. Since we hope that FPS will fill in a gap with existing web platform solutions in preparation for third-party cookie deprecation, we expect to learn from them how FPS is leveraged by site authors. As the list of sets grows over time, and the ecosystem adapts to a post-third-party cookie world, we can also mature the process to the point where we can consider alternative decentralized schemes such as the one proposed. With the current process, we expect to institute set lifetimes, which will

		allow us to evolve the intake process over time. We can revisit this idea once the submission process matures.
Repository moderation	Enact community moderation of the FPS Submission repository to prevent abuse. Bad actors can easily overwhelm the process employing burner origins to propose sets, and an overwhelming volume of requests might affect operations for a genuine set proposals.	We are trying to make the checks as objective as possible by relying on technical validation checks. We think this is the most scalable approach to the submission process. In keeping with this goal, we will also aim to ensure the process is resilient to spam/burner submissions.
Associated Subsets	Will FPS be able to support third-party Vendor/SaaS flow use cases through Associated subsets?	The third-party vendor/SaaS flows are not a use case that is currently considered in scope for FPS. We welcome additional feedback on how cross-site cookies are used for these use cases here.
FPS + CHIPS integration	Request for FPS + CHIPS integration in order to support use cases such as A/B testing.	We are discussing this use case and are also considering discussing this further in a WICG call and welcome additional input here
GDPR	Proposal for a new FPS subset to be modeled after GDPR concepts.	We discussed this proposal internally and weighed it against other feedback received as well as our privacy goals. We've provided an answer explaining why we will not be pursuing this proposal at this time.
Memory	Expected change in browser memory size when the FPS list is incorporated	There have been precedents for browsers to store these kinds of lists with minimal memory impact, such as the Disconnect Tracking Protection List. While the FPS list will be copied to each Chrome client locally, we will continue to monitor the file size and are confident that we can optimize the memory footprint.

Fenced Frames API

Feedback Theme	**Summary**	**Chrome Response**
Fenced Frames limitations	Clarity around the limitations imposed by Fenced Frames	In March, we updated our explainer on Fenced Frames which provides information on its capabilities and we welcome any additional feedback.
Expand access information	Request to expand access to information around neighboring frames	We are looking to further understand why this is a requirement from the ecosystem, and we welcome any additional feedback.

Fenced Frames and iFrames	Questions regarding the feature parity between Fenced Frames and iFrames	All available Privacy Sandbox APIs and reports will be available for iFrames and for FencedFrames in the same way.
Re-sizing Fenced Frames	Restricting frame size changes affects certain use cases.	We are interested in learning more about the types of use cases that are affected by the restriction and would welcome additional feedback here.

Shared Storage API

Feedback Theme	**Summary**	**Chrome Response**
Third-party Worklets	Can third parties write to Shared Storage, partitioned by origin? Or call other worklets for third-party measurement?	The browsing context's origin of where the code is being executed determines whose shared storage that data is written to. When a third-party code is added to a page, the third-party code can be embedded as an iFrame with its own browsing context which allows the third-party code to write to its own origin. The third-party code can also be embedded as a script instead of an iFrame, which does not switch the browsing context, and the third-party can write to the embedder's shared storage. Note that only the owner of that shared storage can read from that shared storage.
Deduplication	Deduplication would not be possible for interactions outside the Chrome ecosystem.	Shared Storage is meant to provide Chrome browser based unique reach outputs within Chrome. We are interested in working with ad techs to understand how these outputs can be used as a part of their broader reach models. We understand that the outputs themselves might only account for a portion of interactions and are interested in working with adtechs to explore additional modeling methodologies that could be layered on top.
Conversion Look back Window	Request to have lookback window for conversion rate in order to see changes in conversion over time.	This ask can be implemented by processing the various conversion paths on the client-side using Shared Storage which affords additional flexibility for advanced analytics over secure unpartitioned browser storage.
Item Expiry Window	Request to extend the expiry window to 90 days.	The data retention policy was updated in November 2022, and states that each key is cleared after thirty days of last write. We

		welcome additional feedback to understand if the new policy will work for the ecosystem.
Creative Rotation	Creative rotation use cases do not reflect actual actions post-auction	We are interested in hearing from more buy side ad tech companies on whether the Creative rotation documentation is accurate or not.

CHIPS

No feedback received this quarter.

FedCM

Feedback Theme	Summary	Chrome Response
Identity assertion endpoint	Explicitly allow arbitrary request to the identity assertion endpoint.	We have been collaborating with Mozilla on this pull request to limit websites' ability to make cross-origin credentialed requests silently without causing user annoyance; and will continue reviewing and addressing other feedback as well.
Pre-populate identity	Can FedCM be used to pre-populate a sign-in form with an identity provider from the FedCM list?	The concern for this use case is that it may result in the leaking of information when a site that has not engaged with the user is able to query the last IDP used by the user. We are discussing this issue further here and welcome additional feedback.
Contextual Account Selection	Proposal to add contextual signals in the account selection UI.	We are considering this proposal and welcome additional discussions here.

Fight spam and fraud

Private State Token API (and other APIs)

Feedback Theme	Summary	Chrome Response
Capabilities Gathering Survey	In early Q1, we finished collecting our survey results for which capabilities are needed for various anti-fraud use cases, and shared	We plan to incorporate this feedback as we develop new proposals and prototypes for purpose built, privacy preserving APIs for anti-fraud capabilities. We expect we will

	them publicly (minutes, results).	prioritize development where there is sufficient need and there is existing technology we can build upon to introduce the capability to the web, while preserving user privacy. For example, device and boot integrity was highly ranked and many platforms have existing APIs that securely share an assessment of device integrity, so it is a good candidate to pursue exploration within the community groups.
PST Intent to Ship Feedback	As part of the intent to ship, we received a concern in proceeding given that we are utilizing an older version of Privacy Pass. We also received feedback that the specification was unclear in certain sections, and should be improved to facilitate browser compatibility.	We plan to implement many of the suggested specification changes before shipping to GA, as well as a few API changes. The feedback came right at the end of Q1, so we are following up on the GitHub issues with specific details and an update to our launch plan (in progress, as of the publication of this report). For the larger changes to the API, we are open to considering them, but we feel the best way forward is to proceed with launch to GA and get hands on feedback from more developers. We hope to continue this discussion and pursue browser standardization. If, and when, a new standard would emerge, we will consider adopting and developing a plan to carefully transition to it.

Google Roadmap for Effectiveness Testing of the Privacy Sandbox Proposals

As we continue to approach the deprecation of third-party cookies, efforts to invest in testing the effectiveness of the APIs are increasingly becoming a priority. The Privacy Sandbox Team is planning to publish in consultation with the CMA a proposal for facilitating tests of the Privacy Sandbox proposals during the course of May 2023.

For its part, Google Ads is beginning to undertake initial testing to road test the APIs and provide feedback to the CMA and the ecosystem. Google is conscious of the importance of transparency for the ecosystem, so that they can plan their investments and forecast participation in future tests, and as such has included Google Ads' testing plans below:

Topics API for Interest-based Advertising:

- During Q1 2023, Google Ads completed running an Interest-based Advertising experiment on Origin Trial Chrome Desktop + Mobile Web traffic, utilizing a combination of privacy-preserving signals including contextual information, the Topics API from the Privacy Sandbox and first-party identifiers such as Publisher Provided IDs.
- In consultation with the CMA, Google Ads published a white paper that outlines the methodology and shares the results of this experiment.

Measurement APIs:

- In Q2, Google Ads envisages publishing API integration guidance for third-party ad tech on how the Event and Aggregate-level APIs could be combined
- In Q3, Google Ads envisages publishing guidance on how ad tech could improve the performance of the Event and Aggregate-level APIs via API-based mitigations.

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

Updates on User-Agent Reduction

Rollout of User-Agent Reduction

During this reporting period Google has provided the CMA and the ecosystem with information regarding its efforts to limit passively shared browser data through User-Agent Reduction ('UAR'). In an effort to increase transparency, Google has coordinated with the CMA to publish these updates.

In particular, as announced in the blink-dev email thread, Google started gradually rolling-out UAR during Q1 2023. While progressing with the implementation of UAR, in line with Google's cautious approach regarding latency and ecosystem dynamics, we initially

limited the March 21 increase of Phase 6 User Agent Reduction to 5%, rather than 10%. As we move forward with the implementation, we will continue updating the ecosystem and the CMA on any relevant developments, particularly in cases where there would be a justified reason to delay rollout compared to the envisaged timeline.

The current envisaged timeline for the roll-out of UAR Phase 6 is as follows:

Stable 1% [Completed]: February 21, 2023

Stable 5% [Completed]: March 21, 2023

Stable 10% [Completed]: April 4, 2023

Stable 50%: April 25, 2023

Stable 100%: May 9, 2023

This updated timeline and all other timeline updates can be found on the blink-dev email thread.

Latency impact measurement

As part of the broader evaluation of the gradual rollout of UAR, Google has been asked by the CMA to measure certain aspects of the latency impact on the ecosystem. We investigated whether sites that use the Critical-CH response header incurred a meaningful latency impact by measuring and comparing the First Contentful Paint (FCP) metric against a local build of Chrome with the feature enabled and another with the feature disabled.

We tested 60 sites (randomly selected) with an automation framework to load the sites hundreds of times. In aggregate, the first page load across all sites appeared to incur an additional ~100ms in its FCP. This aggregate figure should however be read with caution, as it is not representative of real usage, each site's architecture will have a different influence on latency, and, because of this great variance, the aggregate analysis is not necessarily meaningful. We think that a more representative figure can be established when looking at a site in isolation. When taking this perspective, the latency impact figure is cut in half to ~50ms. On any subsequent navigation or visit, the FCP delta was merely ~3.5ms. Based on this result, we suspect this figure is more representative of real usage, especially for users across the globe visiting sites making use of edge caching or other CDN strategies.

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 the underlying thinking. Google continues to respond to a continuous sequence of detailed questions in this respect. As part of this, Google and the CMA continue to operate a joint process by which the CMA carefully reviews relevant Google announcements before they are published.

We engaged closely with the CMA and with a variety of representative stakeholders across the industry as part of the process of updating the Topics taxonomy, which we will announce in Q2. The results of our engagement has so far given us positive expectations about the improved utility for the ecosystem, but we look forward to hearing more comments on the revised taxonomy once it will be published, and we'll continue engaging with feedback throughout the upcoming months. Both Google and the CMA continue to ensure that design updates like these are in line with the Commitments.

CMA concerns

The CMA has not during the relevant period expressed concerns for resolution pursuant to paragraph 17(a)(ii), nor notified any such concerns pursuant to paragraph 17(a)(iii) of the Commitments. However, 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:

First-Party Sets - The CMA has shared with Google that some stakeholders believe that there is ambiguity in Google's communications regarding FPS. To clarify, Google's intent is to have three associated domains in addition to the primary domain (four total). This requirement, and further details on the submissions process, can be found in the Submissions Guidelines and the Chrome Developer Blog. Google will continue to work with the CMA and the ecosystem to improve the readability and clarity of these resources.

Topics - In Q2 2023 Google will announce some important updates to the Topics API, particularly the publication of an improved taxonomy compared to the one originally launched for testing. The taxonomy is the list of available topics that may be returned by the API. We repeatedly received feedback that the testing taxonomy did not represent the topics that the advertising industry cared most about, and that's why we are working to add some important categories and remove ones with lower utility, reflecting Google's own internal assessment as well as feedback from the ecosystem.

One of the key issues that we have considered as we worked on the revised taxonomy is the granularity, which was also the focus of stakeholder feedback in the past quarter. The CMA shared that some ecosystem participants feel that the taxonomy should not be made any more granular due to privacy concerns, whereas others conversely would encourage more granularity. We are conscious of the importance of these various considerations, and we look forward to receiving feedback on the improved taxonomy once it will be published, which will hopefully contribute to addressing the views expressed by both groups of stakeholders.

User-Agent Reduction - The CMA has continued to highlight concerns related to UAR, and in particular latency metrics. Google is in regular communication with the CMA concerning its plans for the Phase 6 Rollout. See the section above for further details.

Standards Development - The CMA has shared that some stakeholders expressed the view that W3C is an unfamiliar venue for publishers, and is not able to facilitate non-technical "policy" discussions. This wider issue concerning requests for additional forums for non-technical ecosystem players has been addressed in the feedback tables under General Feedback.

Timeline - The CMA shared that some stakeholders are still uncertain as to whether Google will meet the announced timeline for the phasing out of third-party cookies. Google is committed to third-party cookie deprecation and is investing significant time and resources into the APIs to ensure they meet the ecosystem's expectations as effective alternatives to third-party cookies while satisfying the Development and Implementation Criteria set out in the Commitments. The development of the Privacy Sandbox APIs is progressing at pace. The Privacy Sandbox APIs are already available in OT for testing and will be generally available for 100% of Chrome traffic this summer.

Alternatives - The CMA has shared that some stakeholders are keen to ensure that Google's technologies do not close off legitimate alternatives to Privacy Sandbox. Google's efforts are focused on developing the Privacy Sandbox proposals in such a way that they comply with the Development and Implementation Criteria set out in the Commitments, and achieve the purpose of protecting privacy while replacing use cases critical to a thriving web ecosystem. Google welcomes efforts to develop alternative privacy-preserving technologies to support ads targeting and measurement. While

encouraging the development and testing of such technologies, Google will always keep in mind the privacy, safety, and security of its users.[2]

Status Meetings

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

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

Standstill

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

Compliance statement

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

[2] See Google's Q2 2022 Progress Report, page 22.

[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 31 March 2023, Google has complied in the preceding three-calendar-month period with the obligations relating to:

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

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

Signed.....[redacted].....

Full name.....[redacted].....

Date.....[redacted].....

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