



Privacy Sandbox Progress Report

Q4 Reporting Period – October to December 2022

Prepared for the CMA, 25 January 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 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 Q4 2022 updates, covering the period from October 1 to December 31, 2022.

Google is working toward the removal of third-party cookies in H2 2024. In order to get there, Google is taking a phased approach:

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.

Discussion: Already completed, the technologies and their prototypes are discussed in forums such as GitHub or W3C groups. Some limited testing of solutions might happen at this stage to facilitate discussions.

Pre-Launch Testing: Currently ongoing, the technologies for the use cases* are available for testing via Chrome Origin Trials or other pre-launch methods. Changes may be made based on testing results and ecosystem feedback.

General Availability: From Q3 2023 onwards, the technologies for the use cases* are launched and available for ~100% of Chrome traffic. Chrome expects refinements and optimizations as more companies test and use the APIs over time.

Third-party cookie phase out: In Q3 2024, Chrome will phase out support for third-party cookies over a two-month period.

*The use cases are: 1) Fight Spam and Fraud on the Web; 2) Show Relevant Content and Ads; 3) Measure Digital Ads; and, 4) Strengthen Cross-site Privacy Boundaries

Privacy Sandbox Q4 2022 Timeline Updates	
October Timeline Updates	<ul style="list-style-type: none">Changed the name of “Trust Tokens” to “Private State Tokens”²
November Timeline Updates	<ul style="list-style-type: none">Added “OT Closed” to Federated Credential Management
December Timeline Updates	<ul style="list-style-type: none">Added “Feature Flag” to First-Party SetsAdded “OT Closed” to CHIPS API

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

² On October 17, 2022 Google informed the CMA that it was changing the name of Trust Token API to Private State Token API. Google also updated the [privacysandbox.com](#) website to reflect this change. A GitHub explainer for Private State Token API is available [here](#).

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 standard 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
<p>(Also reported in Q3)</p> <p>Usefulness for different types of stakeholders</p>	<p>Concerns that the Privacy Sandbox technologies favor larger developers and that niche (smaller) sites contribute more than generic (larger) sites.</p>	<p>Our response is unchanged from Q3:</p> <p><i>“Google has committed to the CMA 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 these commitments.</i></p> <p><i>As testing of the Privacy Sandbox progresses, one of the key questions we will assess is how the new technologies perform for different types of stakeholders. Feedback is critical in this respect, especially specific and actionable feedback that can help us further improve the technical designs.</i></p> <p><i>We have worked with the CMA to develop our approach to quantitative testing, and are supportive of the CMA publishing a note on experiment design to provide more information to market participants and an opportunity to comment on the proposed approaches.”</i></p>
<p>(Also reported in Q3)</p> <p>Documentation requests</p>	<p>Requests for more resources detailing how to manage testing, analysis, and implementation.</p>	<p>Q4 Update:</p> <p>We appreciate that developers have found our current material helpful, and continue to be committed to providing more material so developers can understand how the new technologies can work for them. Over the past quarter, we added a "News & Updates" section to privacysandbox.com and published</p>

		<p>an extensive review of how the Privacy Sandbox can help drive ad relevance in the future.</p> <p>We've also held public developer office hours sessions to share best practices and demos, along with Q&A sessions with product and engineering leads to allow for live discussion/questions.</p>
Core Web Vitals	How does Privacy Sandbox API latency impact Core Web Vitals?	<p>Keeping latency to a minimum is a key design goal of the Privacy Sandbox APIs. Our current expectation is that API latency should have minimal impact on a site's Core Web Vitals, as the majority of APIs are called after the initial rendering of the website. We continue to monitor and make improvements to reduce latency further for each of the APIs, and encourage continued testing and feedback.</p> <p>Latency in the real time bidding process is addressed in the FLEDGE section under "Performance of FLEDGE Auctions".</p>
Interoperability	Concerns regarding interoperability with other potential solutions	<p>The goal of Privacy Sandbox is to protect users against cross-site tracking while supporting the needs of the web ecosystem. We seek to accomplish this by moving away from legacy browser technologies that enable such cross-site tracking, like third-party cookies, and providing in their place new technologies purpose-built to support specific use cases.</p> <p>The Privacy Sandbox proposals improve privacy by limiting the data that leaves a user's device. The proposals do not place technical restrictions on a website's ability to share or otherwise process data once collected from the browser. The technologies therefore do not prevent companies from entering into "data stewardship" agreements or</p>

		<p>any other similar contractual relationship. Likewise, they do not restrict the ability of users to consent to sharing their data via other means.</p> <p>For clarity, Google has committed to apply the Privacy Sandbox technologies in the same way to all websites, including Google products and services. After Chrome ends support for third-party cookies, the commitments also make clear that Google will not use other personal data, such as users' synced Chrome browsing history, to track users for the targeting or measurement of digital advertising.</p>
--	--	--

Show Relevant Content & Ads

Topics

Feedback Theme	Summary	Chrome Response
Impact on Google search ranking	Enquiry on whether a website's Topics API support will be used as a potential signal for Google Search results ranking.	Some websites may choose to opt-out of the Topics API. The Privacy Sandbox team has not coordinated or requested from the Search organization that they use page ranking as an incentive for websites to adopt the Topics API. Google has confirmed to the CMA that Google Search will not use a site's decision to opt-out from the Topics API as a ranking signal.

Topics classifier	Add url and page content in addition to hostname to determine a webpage's Topic in order to improve utility for various stakeholders.	<p>A user's browsing history is currently classified using a website's hostnames. Chrome continues to evaluate options for considering page level metadata (such as all or some components of the page URL and/or content) in Topics classification. Any utility improvements must be weighed against the privacy and abuse risks.</p> <p>For example, with respect to metadata in particular, a few of the risks include:</p> <ul style="list-style-type: none"> - Sites modifying page-level metadata as a method to encode different (and potentially sensitive) meanings into topics; - Sites modifying page-level metadata to misrepresent their topics for financial gain; - Sites modifying page-level metadata dynamically as a method of cross-site tracking
(Also reported in Q3) Impact on first-party signals	Topics signal may be highly valuable and as a result devalues other first-party interest-based signals.	<p>Our response is unchanged from Q3:</p> <p><i>"We believe interest-based advertising is an important use case for the web, and Topics is designed to support that use case. As described [in our Q3 report], other ecosystem stakeholders have expressed concerns that Topics may not be useful enough to provide value. In all cases, improvements to the taxonomy are an ongoing effort, and we expect the taxonomy to evolve with ecosystem testing and input."</i></p>
Updating Taxonomy	How will the taxonomy list be updated?	<p>We are actively seeking feedback on the taxonomy that would be most useful for the ecosystem. The taxonomy included in the initial Topics API proposal was designed to enable functional testing. Chrome is actively considering multiple approaches for updating the taxonomy. For example, Chrome may utilize a notion of commercial value for each topic to determine which category to include in future iterations.</p>

Topics regional classifier performance	Topics classifier performing poorly in regional domains.	<p>Improvement to the classifier is an ongoing effort. Based on the feedback we have received, one possibility under consideration is to expand the Topics override list, which our analysis shows will increase global coverage and improve accuracy.</p> <p>To explain, the Topics API classification has two relevant components: (1) An override list containing the top 10k sites and their topics, and (2) an on-device ML model that classifies hostnames into topics. By expanding the override list (1), we can improve performance of classification for those regions in which the classifier may be performing poorly.</p>
One week epoch	The one week epoch ³ is too long for users looking to make shorter term decisions.	We are actively looking at what the suitable length of epoch should be and we welcome further feedback on what would be a better epoch for the ecosystem.
HTTP header retrieval	Concern that there is not enough information regarding the HTTP header retrieval of topics.	Work is in progress for headers and fetch(). We also have information available here . We have also added skipObservation information to the explainer.
Topics only aims to help advertisers, not users	Topics/Privacy Sandbox appears to be an industry focused approach. Benefit for users is not as clear as benefit to industry.	We believe the benefit to users is that Topics supports interest-based ads that keep the web free and open, and we also believe it significantly improves privacy compared to third-party cookies. Removing third-party cookies without viable alternatives may negatively impact publishers, and could lead to worse approaches ⁴ which are less private, are not transparent, and are not realistically resettable or controlled by users. Many companies are actively testing Topics and Sandbox APIs, and

³ An epoch is the time period which the Topics API considers when calculating the Top Topics. Currently, an epoch is set to 7 days. That means, on every 8th day, the Topics API looks back at days 1-7 to calculate a user's Top Topics.

⁴ See this [post](#) on GitHub detailing research on the deployment of less private approaches in response to cookie blocking.

		<p>we're committed to providing the tools to advance privacy and support the web.</p> <p>The W3C Technical Architecture Group has recently published its initial view about the Topics API, which we will be responding to publicly. At this stage, since Google has received questions from the ecosystem about what this review may imply for the development and launch of the Topics API, we would like to reaffirm our plan to make it available in Chrome Stable this year. While Google appreciates the input of the W3C Technical Architecture Group, it considers it of paramount importance to continue the efforts to develop and test Topics in consultation with the CMA and the ecosystem.</p>
Data leakage	Concern that Topics may be leaked to other sites without permission.	The design of the Topics API makes it quite unlikely that data from a single publisher (and even a smaller group of publishers) can be leaked in any way. Publisher websites are also in full control over the Topics API and they can prohibit access to this API via permission policy.
Lack of advertisers for testing	Publishers are concerned that they currently are unable to demonstrate the value of Topics to advertisers	In the second half of 2023, we plan to have all the ads related APIs available for integration testing and enable ecosystem analysis of the value of Topics for advertisers. Testing and publication of the results will be supervised by the CMA, which will review the data, analysis and methodology. The ecosystem is encouraged to share feedback with Google and the CMA.
Topics and FLEDGE	Request for more information on how to use Topics within FLEDGE's bidding logic	It is possible to use Topics within FLEDGE's bidding logic. An integration guide is also in progress, and will include additional details on implementation.
Custom ranking for topics caller	Allow rankings to be tailored by caller	The challenge with enabling per-adtech topic ranking/values is that it is a mechanism by which they

		can influence the Topics that are returned, and therefore a fingerprinting vector.
Topics caller priority list	Allow callers to provide a ranked priority list of topics that the Topics API will return based on eligibility.	We are currently discussing this idea further and welcome additional inputs.

FLEDGE

Feedback Theme	Summary	Chrome Response
Google Ad Manager	Concern that Google Ad Manager is the end decider for FLEDGE auctions and would favor Google Publisher Tags and Google Ad Manager.	FLEDGE allows each publisher to choose the structure of the auction, including the choice of top-level and component sellers. Each buyer and seller in a component auction knows who the top-level seller is, and can choose whether or not to bid.
Not enough participants testing FLEDGE	Request to encourage more companies to test FLEDGE, for example by improving the API's functionality and discouraging privacy-intrusive alternatives like fingerprinting	<p>The Privacy Sandbox is proceeding in stages, in close partnership with the guidance of the CMA and ICO, and functional FLEDGE testing has demonstrated necessary stability and capability. Google continues to encourage the ecosystem to test the Sandbox APIs, recently publishing its "Maximize Ad Relevance" documentation to showcase how FLEDGE and other APIs can help support critical use cases for the ad industry after third-party cookie deprecation.</p> <p>Other parts of the Privacy Sandbox already support mitigations to cover tracking (see UA-CH, IP Protection, and Bounce Tracking Mitigations) and will continue to improve over time. Google's goal is not to make FLEDGE the only viable targeting solution, but instead remains committed to working in partnership with industry and regulators to drive the best privacy-preserving ad technologies in the Chrome browser.</p>

Machine learning use cases	More guidance on how machine learning use cases to train auction bidding algorithms will be supported in FLEDGE and Attribution Reporting.	We recognize the need to help testers find the most useful ways of applying the Privacy Sandbox technologies. We have begun to publish guidance specifically related to the use of the various aspects of the Privacy Sandbox APIs as inputs to machine learning. The most recent piece, “Maximize Ad Relevance” , discusses how the ads industry can leverage these signals for machine learning, and we plan to continue publishing such guidance going forward.
Querying FLEDGE Key Value (K/V) Server	Why is the K/V server publicly queryable?	The K/V server aims to provide real-time signals to FLEDGE auctions. As such, the K/V server needs to be accessible from where those FLEDGE auctions execute, which is on user devices, requiring that it be publicly available. A value stored in the K/V server can only be retrieved by a party that already has its key — so if an ad tech only gives the keys to browsers that are in the Interest Group, and does not use keys that can be randomly guessed, then only browsers that need the Value to run their auction will be able to retrieve it.
How to do date/time targeting?	Support for date objects in the bidding logic function.	There are multiple ways to do this. Buyers can ask their seller to provide the current date and time, and it should be easy for sellers to provide this information to all buyers. Buyers can also provide the date and time in their real time key-value response. Finally, buyers can provide the date and time as part of their contextual response in the per-buyer-signals , which a seller could pass to the buyer's generateBid script.
User preferences	Ability for users to choose to block creatives by advertiser when served via FLEDGE, or alternative solutions.	Users have the ability to opt out of Ads APIs in Chrome. 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.
Clearer timelines	Request for more information on availability of privacy protections	We plan to publish more detailed timelines in Q1.

	in FLEDGE, such as requiring Fenced Frames.	
Reporting confusion	Request for more clarity on how FLEDGE reporting will work with other APIs such as Fenced Frames and Private Aggregation API.	We plan to publish an explainer on the interaction between Private Aggregation API, FLEDGE, and Fenced Frames in the coming weeks.
Real-time bidding and FLEDGE	Guidance on how FLEDGE integrates with standard real-time bidding.	The two main things that complicate an ad-tech's ability to do real-time bidding are access to event level data and easier integration into the Attribution Reporting API. We are planning on sending updates and explainers on both of these in Q1.
Performance of FLEDGE Auctions	Report from testers that FLEDGE auctions have high latency	<p>We appreciate reports from testers sharing their results and use cases and have shared some suggestions on how to improve the performance of FLEDGE.</p> <p>In parallel, we have added tooling to the browser allowing developers to better diagnose what is making auctions slow, and have been systematically addressing the primary sources of latency observed. Recent improvements include timeouts for slow auctions, a fast bidder filtering technique, a way to reuse FLEDGE worklets to avoid paying startup costs, and ongoing work to allow the contextual ad request to run in parallel with the FLEDGE startup time and network fetches. We expect latency optimization to continue as an ongoing conversation between Chrome developers and FLEDGE testers based on their real-world experience using the API.</p>
Interest Group size memory limit	Request to raise the limit on the size of a single interest group from 50kB.	We are actively considering the request and are looking for feedback on what limit value works .
Combining the FLEDGE served data with first-party cookie	Will FLEDGE support integration with an advertiser's first-party data?	FLEDGE was built to support advertising using the first-party data an advertiser already has. However, FLEDGE does not intend to support an advertiser learning a person's browsing behavior on any website other than the advertiser's own site.

		<p>Attaching off-site browsing behavior to first-party data is contrary to the goals of Privacy Sandbox.</p> <p>We are planning to share integration guides with more details on how FLEDGE will support integration with first-party data in the coming weeks.</p>
K-anonymity value	How will the value "K" to "k-anon" be decided and will it be published?	The "K" value is still being finalized and we will share more information as our plans develop. We are interested in learning more about how an unknown k value may hinder FLEDGE preparedness and scoping ML model training and we welcome additional feedback on this subject.
Supporting multiple SSPs	How will multiple SSPs be supported in FLEDGE	FLEDGE supports multi-seller auctions as noted in this proposal .
Visibility of bidding logic	Concern that DSP bidding logic will be exposed in JavaScript	With the current design bidding logic JavaScript can be accessed by others, but we welcome more feedback as to why this could be a source of concern for DSPs.
prebid.js	What is the timeline in supporting prebid.js in FLEDGE	Only versions 7.14 and later of Prebid.js support the FLEDGE module. Any publishers interested in testing must add the FLEDGE module and upgrade their Prebid instance.
User defined functions in FLEDGE	How will user defined functions (UDF) be supported in FLEDGE? These are functions that can be programmed by end users to extend the functionality of the API	Explainer available here . This is still being fleshed out so we welcome additional feedback on use cases.
Relaxing same-origin constraint on Interest Group resources	Request to relax same-origin constraint on Interest Group resources to enable certain ad tech use cases	<p>In the current implementation of FLEDGE, <code>biddingLogicUrl</code>, <code>biddingWasmHelperUrl</code>, <code>dailyUpdateUrl</code> and <code>trustedBiddingSignalsUrl</code> must have the same origin as the Interest Group owner.</p> <p>The constraint exists to prevent certain exploits by attackers, as explained here.</p>
interestGroup Ownership	Request to limit whether an ad tech can use <code>joinInterestGroup</code> for the same Interest Groups across sites	Our focus is on how audiences are used, not how they are built. We are discussing potential approaches here and welcome additional input.

Key Value Server Key Expiration	Discussion on removing server keys once the corresponding interest groups have expired	We are exploring ways to handle key expiration and are looking for feedback here .
---------------------------------	--	--

Measuring Digital Ads

Attribution Reporting (and other APIs)

Feedback Theme	Summary	Chrome Response
Origin Trial traffic	Current Origin Trial traffic is not enough to conduct utility testing.	The current Origin Trials are meant for ecosystem players to conduct functional testing in order to ensure the API is working as intended. We understand that larger amounts of traffic will be required to perform utility testing once the development of the various Privacy Sandbox API is more mature. The current testing timeline envisages that this will occur by General Availability (i.e. when the technologies for the use cases will be launched and available for 100% of Chrome traffic) at Q3 2023 (see our up-to-date timeline on privacysandbox.com). We welcome any additional feedback on use case testing that requires additional traffic.
Overlap in functionality of different Privacy Sandbox measurement APIs	Concerns in having multiple measurement approaches overlap through Privacy Sandbox will increase complexity, for example, Attribution Reporting API and Private Aggregation API.	We are working on better documentation to clarify the different use cases for the APIs, and welcome additional feedback on what areas are lacking explanation. For example, Attribution Reporting API is intended specifically to support conversion measurement, whereas Private Aggregation API and Shared Storage are general-purpose APIs intended to support a broader range of cross-site measurement use-cases.
Retry failed report request	Clarification on how many times a report request is attempted if it fails.	We have published guidance on this . To summarize, reports are only sent when the browser is running/online. After the first failure to send, the report is retried after 5 minutes. After the second failure, the report is

		retried after 15 minutes. After that, the report is not sent.
Reporting Delay	What is the expected reporting delay?	We are looking to hear more feedback from the ecosystem on any reporting delays they are experiencing as we collect data to further assess these delays in parallel.
Prerender pages	Will Attribution Reporting API attribution work on prerender pages?	Attribution registration is deferred on prerender pages until activation (actual click or view takes place). This means we will defer the `attributionsrc` request ping.
Measuring conversion lift	How to measure conversion lift with A/B testing on the same domain.	Websites can measure conversion lift with A/B testing on the same domain through attribution reporting. They can encode their A/B parameters as keys using the aggregate API and then receive summary reports for the conversion values by those key buckets.
(Also reported in Q3) Cross-domain conversions	How to track the conversions that are cross domain, such as with 2 or more destinations?	Q4 Update: We have published a proposal to remove the landing page destination restriction which enables cross domain conversions to be tracked. This proposal has been implemented.
(Also reported in Q3) Expiry setting in conversion report	Request to support report filter / expiry for less than 24 hours.	Q4 Update: We have shared this pull request which will decouple expiry and reporting windows to mitigate the trade off of reporting delay vs conversion expiry. This is now launched in M110.
Fraud and Abuse	Requests from advertisers and marketers to be able to slice and aggregate data based on publisher sites where their ads are served, which would also give more insight into potential fraudulent ad practices.	This feedback is actively being discussed here and we welcome additional inputs.
(Also reported in Q3) Event level reporting delay	The delay of 2-30 days in event level reporting may be too long for certain use cases.	Event level reporting has default reporting windows of 2, 7, and 30 days. This can be modified by using the "expiry" parameter. Ad-techs could configure the expiry, with a

		<p>minimum of 1 day, to get potential reports in less than 2 days. We limit the granularity of expiries to 1 day as a privacy protection mechanism, as more fine-grained reporting could result in timing attacks. Additionally, we allow setting independent "expiry" parameters for event level and aggregate reports. See here.</p> <p>Additionally, Google Ads will not get any special reporting windows that other ad-techs do not get via the Attribution Reporting API.</p>
Same reporting origin requirement	Request to remove requirement for source registration origin to be the same as the conversion registration origin.	We propose using HTTP redirects to delegate registration to solve this use-case. We welcome any additional feedback on the new guidance.
Conversion tracking	Need to differentiate whether the conversion happened before/after certain hours set by the advertiser.	Attribution Reporting API supports setting an expiry window and priority for source attribution. By using both, it will technically be possible to attribute a conversion that happened within X days window separately from a conversion that happened after X.
Noise simulation	Request to be able to simulate the different volume of conversions per bucket, to understand the impact on advertisers with less conversions	We are looking to add ways to simulate this in future versions of Noise Lab . We welcome any additional feedback.
Reporting on mobile	Will the report still be sent when Chrome is running in the background on mobile?	At the moment, even on mobile, the report will not be sent when Chrome is in the background. This is likely to change when the API integrates with Android Privacy Sandbox. See here . Note that Android Privacy Sandbox is not part of the Commitments accepted by the CMA.
Data availability	Concerns that Google will have additional access to data via Privacy Sandbox APIs	<p>First, Google Ads does not receive any preferential access to data from the Attribution Reporting API or other Privacy Sandbox APIs. This issue is also addressed in the General Feedback section under "Interoperability" which includes more detail on Google's Commitments.</p> <p>Second, on the difference between larger and smaller sites, Google</p>

		<p>recognizes that noise-based privacy protections may have a greater impact on smaller data slices. However, there are some possible mitigations: for instance, methods like aggregating over longer periods of time would solve this problem. That said, it remains unclear if conclusions based on very small data slices (like one or two purchases) are meaningful at all to advertisers. During the Origin Trial, Google has encouraged testers to take advantage of the ability to experiment with a wide range of privacy and noise parameters so they can provide more specific feedback on this issue.</p>
--	--	--

Limit Covert Tracking

User-Agent Reduction

Feedback Theme	Summary	Chrome Response
Delay User-Agent Reduction until web ecosystem is more ready	There is not sufficient time to adapt to the coming User-Agent Reduction changes.	We address this feedback under "Stakeholder Concerns" in the section titled "Google's interaction with the CMA".
Delay User-Agent Reduction until web ecosystem is more ready	Request to delay User-Agent Reduction rollout until Structured User Agents (SUA) is deployed.	<p>The Google Ads team proposed the Structured User-Agent addition (see specification) to OpenRTB in October 2021 and it was incorporated in the 2.6 spec update released in April 2022.</p> <p>We have some evidence that SUA is deployed and available today, as demonstrated by the Scientia Mobile blog post demonstrating how to use UA-CH and the WURFL API to create a SUA.</p>

User-Agent Client Hints

Feedback Theme	Summary	Chrome Response
Test UA-CH with other anti-covert tracking techniques	Guidance on how to test all Privacy Sandbox APIs and fingerprinting techniques proposed together in a holistic approach.	<p>Our testing plan was designed in order to reflect the asynchronous timelines for developing some of the anti-fingerprinting measures as opposed to the rest of the Sandbox Proposals. It addresses the reality that some anti-fingerprinting measures (i.e. Privacy Budget, IP Protection, and Bounce Tracking Mitigations) will be fully-developed and ready for launch to General Availability only after third-party cookie deprecation.</p> <p>While those anti-fingerprinting measures will not be included in quantitative tests, they will be subject to qualitative assessment based on the facts available at the time of Standstill.</p>
(Also reported in Q2) Performance	Concerns about the latency of getting hints via Critical-CH (on the first page load).	See dedicated UA-CH section below
Insufficient Feedback	Feedback from the ecosystem about the UA-CH change may not be sufficient, leading to concerns about a lack of awareness from the ecosystem.	<p>We've been proactively sharing our plans to ensure a careful rollout that minimizes disruption.</p> <p>The plans for User-Agent Reduction and the UA-CH API were presented to the W3C Anti-Fraud Community Group on March 18, 2022 and to both the Web Payments Working Group and the Web Payments Security Interest Group on January 20, 2022. No significant concerns were raised during or after the presentations.</p> <p>Google has proactively engaged with more than 100 site operators to obtain feedback. Furthermore, Google has also used Blink-Dev channels to communicate the roll-out of the user-agent reduction publicly based on feedback from ecosystem stakeholders.</p>

Timing	Concerns raised regarding timing of rollout and industry preparedness	See dedicated UA-CH section below
Chrome Platform Status	Requested that the chromestatus page for UA-CH be updated.	The chromestatus entry was updated on December 19 to "Mixed signals".

IP Protection (formerly Gnatcatcher)

Feedback Theme	Summary	Chrome Response
Opt in or Opt Out	Is IP Address Privacy Opt In or Opt Out?	Our goal is to provide IP Protection to all users. With that goal in mind, we are currently evaluating user choice options for IP Protection.
IP Address use case for first-party data	Is it possible to use IP addresses to stitch together user journeys across first-party domains post IP Protection?	As previously published , IP Protection will initially focus on tracking in the third-party context, which means first-party domains will not be impacted.
Ad Tech use cases	How can companies set up anti-fraud measures with IP Protection?	We recognize the importance of IP address as a signal for anti-fraud efforts in today's web. As part of our Commitments to the CMA (paragraph 20), we have said that we will not implement IP Protection without making reasonable efforts to support websites' ability to conduct anti-spam and anti-fraud efforts. One of our top priorities is to understand how IP Protection impacts anti-fraud use cases and detection capabilities, so that we can further invest in privacy preserving technologies that help companies preserve web safety. We encourage feedback and input on new proposals aimed at supporting the needs of security and anti-fraud companies, even as signals change over time.
Fraud and Abuse	Does IP protection include Denial of Service (DoS) Protection?	We are committed to improving privacy while keeping the web safe, and protecting against denial-of-service attacks is an important anti-abuse use case to design for. We hope to minimize impact to DoS protections through both the design of IP Protection itself

		<p>and through new anti-abuse solutions. Because IP Protection is initially focused on third-party embedded services, some stakeholders have indicated that it should have limited impact on DoS protection for first-party sites. However we continue to ask for public feedback to assess risk to DoS use cases, particularly to third-party embedded services.</p> <p>In parallel, we are exploring abuse-feedback and client-blocking mechanisms that would enable a site or service to block a spammy user, without identifying the user.</p>
Content Filtering	Content filtering with IP Protection	<p>Different companies have different needs with respect to filtering content and customizing user experience. Many such use cases do not currently rely on IP addresses and therefore should be unaffected by IP Protection. For example, a publisher looking to tailor its content and drive more engagement might use first-party cookies or third-party partitioned cookies (CHIPs) to understand a user's interests and previous interactions with the publisher. Or an ad tech partner focused on delivering the right ad to the right user can incorporate FLEDGE and Topics, for example, to deliver similar ad outcomes as they do today with third-party cookies or other cross-site tracking technologies.</p> <p>We are also exploring building new privacy-preserving capabilities into IP Protection, such as coarse geolocation, to further support content filtering where existing mechanisms may be insufficient. We welcome additional feedback on content filtering use cases that may be impacted by IP Protection.</p>
(Also reported in Q3) Geolocation use cases	IP Protection may prevent legitimate geolocation use cases from working in the future, such as	<p>Q4 Update:</p> <p>We are working with stakeholders to</p>

	content personalisation based on geolocation.	ensure that Chrome continues to support legitimate use-cases for IP addresses. We are seeking ecosystem feedback on IP Geolocation granularity here .
--	---	---

Privacy Budget

Feedback Theme	Summary	Chrome Response
Clearer documentation	More examples so stakeholders can anticipate how things may be limited once Privacy Budget is implemented.	<p>The Privacy Budget proposal is still under active discussion and has not been implemented by any browsers. The earliest date of scaled availability represents the earliest date when Privacy Budget could be enforced. This will not happen before the removal of third-party cookies in 2024. We do not have any additional documentation to share at the moment.</p> <p>We will share additional details on the proposal when it becomes more finalized. In the meantime, we welcome stakeholders to share any feedback that would help in the development of the proposal.</p>

Strengthen cross-site privacy boundaries

First-Party Sets

Feedback Theme	Summary	Chrome Response
(Also reported in Q3) Domain limit	Request to expand the number of associated domains.	<p>Q4 Update:</p> <p>We have released FPS for local testing, including a mock set submission process on GitHub and a flag to test rSA and rSAFor locally. We have also hosted two open meetings for developers on FPS to continue to address questions around use cases for the associated subset. We encourage developers to test out FPS</p>

		<p>functionality to provide feedback on how the domain limit for the associated subset would impact the usability of FPS for their use cases.</p> <p>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.</p>
Request for more details about the abuse mitigation measures.	What happens if a domain is added to a set they did not consent to?	<p>We have published submission guidelines for First-Party Sets here on December 2, 2022.</p> <p>As explained in the submission guidelines, any set change management will be following and respecting a validation process on GitHub, including validation on ownership, which should mitigate this risk.</p>
Abuse mitigation	Concern that First-Party Set formations can be exploited.	We are looking at ways to expand technical checks for subset types and are actively seeking additional input from the community here .
Ads use cases	Questions on whether First-Party Sets should be used to support Ad targeting	We're not trying to support Ads targeting use cases for First-Party Sets, and we recommend using the Ads APIs available for such use cases.
(Also reported in Q3) Policy	Concern that FPS is not consistent with the CMA Commitments regarding "Applicable Data Protection Legislation", on the basis that GDPR does not impose a limit on the number of sites in a set while FPS envisages a limit of 3.	<p>Our response is unchanged from Q3:</p> <p><i>"Google is continuing to commit to the CMA 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, publishers and advertisers as well as impact on privacy outcomes and compliance with data protection principles as set out in the Applicable Data Protection</i></p>

		<p><i>Legislation. The concern expressed does not disclose any incompatibility with GDPR. We continue to work closely with the CMA to ensure that our work complies with these commitments.”</i></p>
<p>Alternative proposal</p>	<p>GDPR Validated Sets</p>	<p>In addition to the feedback provided by the ecosystem on the proposal to adopt “GDPR Validated Sets,” Chrome has concerns about the following limitations of this alternative proposal:</p> <ul style="list-style-type: none"> - “GDPR Validated Sets” claims to “align to” GDPR (although it is not really clear what is meant by that). In contrast, Google’s commitments require it to take into account “impact on privacy outcomes” more generally. In its decision accepting the commitments the CMA points out that this is distinct from Google’s obligation to take into account “compliance with data protection principles as set out in the Applicable Data Protection Legislation,” which, as the CMA explains, reflects the fact that Google is bound by the Applicable Data Protection Legislation, both as it applies to the Commitments and more generally. - We have privacy concerns about the proposal to allow domains to appear in multiple sets. First-Party Sets are intended to support specific use cases that currently depend on third-party cookies without enabling pervasive cross-site tracking. Allowing domains to join multiple sets would remove a key privacy protection built into the First-Party Sets proposal, without introducing any other meaningful limitations. - GDPR Validated Sets also proposes to “define a Set as a group of data controllers and processors that share a common use policy.” This is similar to the requirement in our original

		<p>First-Party Sets proposal that all parties in a set must share a common privacy policy. We have since removed that requirement based on strong feedback from the ecosystem raising concerns about privacy policy-based requirements. For example, we heard from site publishers that maintaining a common privacy policy was infeasible because of product and geographical variations, among other challenges raised by members of the W3C community (1, 2, 3). We believe that the same challenges would apply to this proposal.</p> <p>Since this alternative was raised, Chrome has updated the First-Party Sets proposal and published submission guidelines for creating new Sets.</p>
--	--	---

Fenced Frames API

Feedback Theme	Summary	Chrome Response
Fenced Frames restrictions during OT	What are the current restrictions around Fenced Frames for the Origin Trial period?	We are working on documentation on the restrictions and implementation status and plan to share it during Q1 2023.
Multiple ads in a single Fenced Frame	Request to display multiple advertisers in one Fenced Frame in one auction	Currently, this request is not being actively developed, but we welcome additional feedback if ecosystem players consider the feature important.
Web Bundles	What are the requirements and support planned for Web Bundles with Fenced Frames?	We currently do not have an update on whether this will be the requirement in the future. Any changes would be announced in advance and would not be enforced before third-party cookie deprecation. Please see this explainer for the current status.

Shared Storage API

Feedback Theme	Summary	Chrome Response
Shared Storage for Ad Tech	Uncertainty surrounding the use of shared storage for Ad Tech use cases.	<p>Shared Storage and Private Aggregation API can be used for different kinds of measurement purposes that need cross-site storage measurement. Some examples are listed here.</p> <p>We are foreseeing DSP and Measurement solution providers to be the main integrator for ads use cases.</p>

CHIPs

Feedback Theme	Summary	Chrome Response
(Also reported in Q3) Partitioned requirement	Add explicit behavior requirement for “Partitioned” attribute on first-party cookies.	<p>Q4 Update:</p> <p>After discussions on GitHub and PrivacyCG calls, the behavior that we have aligned on is that Partitioned cookies set on first-party cookies will use a partition key of (A,A) where “A” is the top-level site. We will document this behavior on the explainer and specification.</p>
Cookie Management	Are there tools for managing/governing first-party or third-party cookies?	<p>Chrome DevTools and NetLog can be used to test sites with third-party cookie blocking enabled. Both tools report when cookies are blocked due to user configuration. We welcome feedback on what sort of additional auditing websites would like to see.</p>

FedCM

Feedback Theme	Summary	Chrome Response
IdP requires knowledge of RP to allow a session	Issue when a user is trying to log into the Feide IdP from two different RPs.	We are discussing potential solutions to this issue here .

<p>Interoperability</p>	<p>Concerns regarding the impact of FedCM on the relationship between users and websites they log into using FedCM, and “interoperability” among websites.</p>	<p>FedCM aims to continue supporting federated-identity services that currently rely on third-party cookies once third-party cookies are removed from Chrome. We expect that FedCM will be just one option available to such services; identity providers (IdPs) and relying parties (RPs) are free to use other technologies that may better suit their needs.</p> <p>It appears that concerns regarding the user-RP relationship and “interoperability” owe to a misunderstanding of the FedCM proposal. FedCM leaves it to IdPs to decide what information to share with an RP, and in what form, once the user has chosen to sign in to that RP’s site. FedCM does not require IdPs to “create a unique pseudonymous identifier for each [RP] with whom the user authenticates.” Rather, FedCM is open for each IdP to choose whether to share the user’s actual identifier, a per-site version of that identifier, or some other version of this information.</p> <p>(The FedCM specification does identify cross-site correlation as a privacy risk associated with the API and discusses directed (per-site) identifiers as a possible mitigation. However, the decision whether to use directed identifiers is left to IdPs, not imposed by the browser.)</p> <p>FedCM also already provides for user choice with respect to identity. For example, if a user has multiple identities with the same IdP (e.g. a work profile and a personal profile), FedCM provides a way for the user to select which one they want to use to log in to the RP’s site. Beyond that, each RP decides for itself which IdPs to support on its site. One aspect of that decision is considering the mechanism that an IdP relies on,</p>
-------------------------	--	---

		whether that's FedCM or a different technology. Again, the browser does not dictate these choices for RPs or IdPs.
--	--	--

Fight spam and fraud

Private State Token API

Feedback Theme	Summary	Chrome Response
Handling Bots	What happens if the issuer discovers that Private State Tokens have been issued to bots?	To avoid tokens issued to bots from remaining in the ecosystem for a long time, issuers should rotate the keys they use to sign tokens regularly so that old tokens issued under potentially broken issuance logic expire and sites redeem newer tokens with updated issuance logic.
Same-site form submissions	Could Private State Tokens be used for same-site form submissions that involve full-page navigation (i.e. Content-Type: application/x-www-form-urlencoded) rather than a request from the fetch/XMLHttpRequest APIs?	This isn't currently supported in the first version of Private State Tokens. We welcome feedback from the ecosystem if there is a strong demand for this use case.
Server-side verification	Questions on whether Private State Tokens can be verified server side?	Tokens are redeemed against the issuer, and then the issuer creates a redemption record that could contain the token itself or some signed value derived from the token, servers can use that redemption record to verify the authenticity of the token, and we expect different issuer ecosystems will come up with different standards for how to interpret their redemption records.

Google Ads 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. 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 for Q1 2023 below:

Testing Topics API:

- During Q4 2022, Google Ads completed an interest-based advertising (IBA) experiment with advertisers, which replaced third-party cookies with simulations of the Topics API (due to insufficient Origin Trial Chrome traffic).
- In Q1 2023, Google Ads is running an experiment utilizing real Topics from Origin Trial traffic on Chrome (Desktop + Mobile Web) for serving Interest-Based Advertising on Display Network inventory available via Google Ad Manager and AdSense. Google is cooperating with the CMA and will aim at publishing the results in coordination with them around the end of Q1 2023.

Testing Measurement APIs:

- During Q1 2023, Google Ads is testing with 5% Origin Trial Chrome Desktop + Mobile Web traffic (from Google Owned and Operated properties). Google is cooperating with the CMA and will aim at publishing the results in coordination with them around the end of Q1 2023. We expect the data to be sparse in Q1 2023.

Beyond Q1 2023, Google Ads currently envisages conducting the following testing of the Privacy Sandbox APIs:

- Q2 2023 -
 - Testing Measurement APIs with 5% Origin Trial Chrome Desktop + Mobile Web traffic (from Display Network inventories).
- Q3 2023 -
 - Testing Measurement APIs with General Availability Chrome Desktop + Mobile Web traffic (from both Google Owned and Operated properties and Display Network inventories).
 - Testing the FLEDGE API on General Availability Chrome Desktop + Mobile Web traffic.

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.

Updates on User-Agent Reduction

During this reporting period Google has provided the CMA with information regarding its efforts to limit passively shared browser data through User-Agent Reduction. In an effort to increase transparency, Google has coordinated with the CMA to publish these updates here, and it will also address some concerns passed on by the CMA.

Rollout of Phase 6: The CMA has received concerns regarding the impact that the rollout of Phase 6 of User-Agent Reduction may have on the ecosystem.

First, Google would like to reassure the ecosystem that all the information currently available in User-Agent strings is recoverable via User-Agent Client Hints.⁵

Second, Google is conscious of the importance of ensuring that the ecosystem is prepared for Phase 6 of Reduction, and it would like to reassure stakeholders that the rollout will occur gradually. Specifically, Google notes that from launch to Stable in Chrome 110 (**Feb 7, 2023**), traffic will be gradually ramped up over several weeks:

- **February 14:** Phase 6 roll-out starts one week after the launch of Chrome 110 Stable with **1% traffic**.
- **March 14:** Roll-out increases to **10% traffic**.
- **March 28:** Roll-out increases to **50% traffic**.
- **April 11:** Roll-out reaches full Stable population with **100% traffic**.

This proposed timeline was posted in the following [Blink intent](#).

In preparation for the rollout of Phase 6, back in November 2022, the Chrome team asked the Partnership and Go To Market teams for their opinion on timeline rollouts. The report from the Partnership and Go To Market teams was overall positive. They gathered feedback from a scaled outreach campaign of 3 waves to 115+ partners across 6 key verticals (CDN, AdTech, CMS, Anti-Fraud, etc) and additional comments from having driven 1:1 engagements of 22 active testers. Additionally, their teams assisted with Chrome presence across 5+ major events (CAB, Wordcamp EU/US, Risk.Ident, W3C Web Payments Working Group & Web Payment Security Interest Group). Their recommendation was to maintain the existing testing timeline based on the lack of negative signals and positive case studies.

Origin Trials: Google has re-opened and extended the [User-Agent Reduction Origin Trial](#) until March 7, 2023 after receiving requests from certain sites to continue testing against the later phases. Google has also [removed traffic limits](#) from the Origin Trial to simplify testing. The Deprecation Trial⁶ [remains open for registrations](#), and is set to expire on May 23, 2023.

⁵ This applies to both desktop and Android mobile User-Agent strings. Note that Android Privacy Sandbox is not part of the Commitments accepted by the CMA.

⁶ A [Deprecation Trial](#) allows a site to opt-in to deprecated or removed functionality for some period of time while they migrate to alternative solutions. The [User-Agent Reduction Deprecation Trial](#) opts a site to receive the unreduced User-Agent string.

Deprecation Trial Registration: Google has seen good levels of registration in its UA-CH Deprecation Trial. Sign-ups have continued to increase, with some expected deceleration as the trial has progressed. Of the 150 origins that have signed up, ignoring any duplicates or invalid entries:

- 43 are categorized as “large” usage sites (~29%)
- 35 are categorized as “medium” usage sites (~23%)
- 72 are categorized as “small” usage sites (~49%)⁷

Costs: Google is aware of the potential cost to developers and has engaged with more than 100 external site operators across a number of verticals, publishing the following developer-centric resources designed to help minimize the cost of adoption:

- <https://developer.chrome.com/docs/privacy-sandbox/user-agent/>
- <https://developer.chrome.com/docs/privacy-sandbox/user-agent/snippets/>
- <https://developer.chrome.com/en/blog/user-agent-reduction-oct-2022-updates/>
- <https://github.com/GoogleChromeLabs/uach-retrofill>

Latency: Some stakeholders have raised concerns around the impact of ad techs having to call information from UA-CH when it is no longer available in the UA string.

Regarding Search, as noted in the [Google Search Central Blog](#), speed is one input of many into search rankings; Google Search’s “Speed Update”, released in July 2018, only affects pages that deliver the slowest experience to users and will only affect a small percentage of queries. As demonstrated below, the median Client Hint fetch latency on Windows (and Android) is about one half of a millisecond. This is not a significant performance regression and therefore should not impact search rankings.

Regarding Ads, internal experiments run by Google Search Ads, Display Ads, and YouTube Ads, indicate that switching to collecting UA-CH had no noticeable impact on Ads metrics. While individual sites will have different performance requirements, Google’s own internal experimental results and deployment experience signal encouraging trends for the wider Ads ecosystem.

Sources of latency can arise from many places: page architecture, choice (or lack) of CDN, the location of a user or its ISP relative to said CDN, optimization (or lack thereof) of page resources, etc. It’s therefore difficult for Google to estimate the exact impact of the work on User-Agent Reduction. Sites that only require low-entropy information from the User-Agent string will not be affected. Some high entropy use cases can be mitigated: third-party embeds are able to obtain the high-entropy client hints via its JS API⁸, and should similarly not observe any material latency increases. Finally, first-party sites that do require high-entropy hints on first-page load have the option to use the Critical-CH header,

⁷ Note that these categorizations are self-reported when a site registers for the deprecation trial, where “small” is from 0 to 10,000 page views a day, “medium” is from 10,000 to 10,000,000 page views a day, and “large” is more than 10,000,000 page views a day.

⁸ The API is `navigator.userAgentData.getHighEntropyValues()` API and this topic was raised in [Issue #134](#).

independently or in conjunction with the HTTP/2 and HTTP/3 ACCEPT_CH frame. This is a one-time cost, as the hints are cached and provided on subsequent visits (until the user clears site data).

All these factors and the data discussed below support Google's view that Rollout of Phase 6 will not have a material impact on third-party websites. That said, to further reassure the ecosystem, Google would like to emphasize that any such potential impact will be addressed and mitigated thanks to the safeguards in place. Google has a monitoring system to alert it to statistically significant regressions in the relevant latency metrics. At this stage, Google has already carried out a number of latency measurements, and it has a budgeted effort in its 2023 plan for performance improvements and latency mitigations if the above processes signal they are needed. As User-Agent Reduction ramps-up, Google will continue to monitor compatibility and latency metrics during the process of reduction, and may pause the rollout or extend the timeline if there is meaningful evidence of a material negative effect. Google will also monitor market impacts, including latency and industry take-up of UA-CH, alongside any technical issues, and take appropriate measures if these metrics raise concerns.

Key metrics and results regarding impact on latency

What are they? Why are they relevant to assess latency?

- **ClientHints.StoreLatency.** This metric measures the latency to store the client hints sent by a given origin in the client's Accept-CH cache. We monitor this to ensure we do not ship unintentional regressions to the underlying client hints system, as this would affect overall page load performance.
- **ClientHints.FetchLatency.** This metric measures the latency to retrieve the client hints for a given origin from the client's Accept-CH cache. We monitor this to ensure we do not ship unintentional regressions to the underlying client hints system, as this would affect overall page load performance and the speed at which the navigator.userAgentData.getHighEntropyHints API can return a result (which we expect to be more common for Third-Party use cases).
- **PageLoad.PaintTiming.NavigationToFirstContentfulPaint.** This metric measures the time from the beginning of a navigation to its [first "contentful" paint](#) (FCP), or the moment when content (i.e., images, text, etc.) is first rendered on the screen. We consider this metric to be a useful proxy for user-observable page load performance.
- **Net.HttpResponseCode.** This metric counts the number of HTTP Response codes encountered. We monitor this to observe if a change results in a statistically significant change in response codes which might affect page load or user experience. For example, an increase in 3XX or 429 codes might indicate an increase in redirects. Other 4XX and 5XX codes, such as 504 Gateway Timeout may indicate errors in how content is sent from the browser, or the server failing to serve a response in a compatible or timely fashion.
- **ClientHints.CriticalCHRestart.** This metric counts the number of HTTP requests restarted because of the Critical-CH response header. This metrics tracks adoption, but an anomalous increase may also indicate a regression causing too many restarts (which are similar to a redirect) which may negatively impact page load performance.
- **Server-side Metrics.** We have requested and received reports from internal and external partners on their own server-side performance metrics, and used this data to prioritize performance engineering investigations and improvements.

Results per metric

ClientHints.StoreLatency

In general, there is no increase to latency for storing Client Hints over the past 90 days:

Android:

p50: 0.508ms steady over 90 days.

p75: 0.762ms steady over 90 days.

Windows:

p50: 0.50ms steady over 90 days.

p75: 0.756ms steady over 90 days.

ClientHints.FetchLatency

Android:

p50: 0.595ms to 0.656ms over 90 days: 0.061ms increase

p75: 0.891ms to 0.985ms over 90 days: 0.094ms increase

A less than a 10th of a millisecond increase does not cause concern today but Google will monitor to make sure this does not steadily increase over time.

Windows:

P50: steady at 0.5ms (hovering between .503ms and 0.506ms) over 90 days

P75: steady at 0.75ms (hovering between .755ms and .758ms) over 90 days

PageLoad.PaintTiming.NavigationToFirstContentfulPaint

Across the board, Google sees improvements to FCP over the past 90 days. This means page loads are faster.

Windows:

P50: 623.537ms to 596.183 over 90 days: 27.345ms decrease

P75: 1257.211ms to 1210.151ms over 90 days: 47.06ms decrease

Android:

P50: 852.403ms to 809.909ms over 90 days: 42.494ms decrease

P75: 1577.316ms to 1496.175ms over 90 days: 81.141ms decrease

Net.HttpResponseCode

Note: everything with a “_20221205” relates to the Stable 5% Phase 5 experiment - [this was ramped up to 10%](#) on January 9, 2023, but it takes approximately 1 week for ramp up before yielding meaningful data.

When comparing the “Enabled_LegacyWindows_20221205” bucket with “Control_LegacyWindows_20221205”, Google does not observe any statistically significant change in HTTP response codes between the groups.

When comparing the “Enabled_NonLegacyWindows_20221205” bucket with “Control_NonLegacyWindows_20221205”, Google does not observe any statistically significant change in HTTP response codes between the groups.

When comparing the “Enabled_DesktopExceptWindows_20221205” bucket with “Control_DesktopExceptWindows_20221205”, Google does not observe any statistically significant change in HTTP response codes between the groups.

ClientHints.CriticalCHRestart

This metric gives a count of requests restarted due to the presence of the Critical-CH response header. It allows Google to track adoption, or observe anomalies that might indicate a regression in functionality.

Over the past 90 days, Google has observed a relatively flat range of being present on 0.0218% to 0.0234% of page loads on Windows, with a slightly positive gradient.

On Android, Google has observed higher numbers: from 0.04% to 0.175% of page loads over the past 90 days.

Server-side Metrics

Google is working with its partner teams, both internal and external, to receive feedback and metrics on client hints latency from the server's perspective. Internal teams have already sent some latency reports which have resulted in actionable feedback and improvements to the client hints implementation. Google has not received new reports related to latency, regressions, or concerns about timelines from these partners.

Community reports

Up to the end of this reporting period, Google has only received two bug reports from the community related to User-Agent Reduction:

Related to Phase 4:

In crbug.com/1317577, a community member reported that some version detection code on their site broke when the BUILD portion of the UA string was reduced to "0" (they were able to fix the broken logic on their site). As a precautionary measure, we [added some code](#) in Chromium that would allow us to freeze the BUILD portion to "9999", controlled remotely via a Finch config. Ultimately we never received any similar reports, so we decided to keep BUILD reduced to "0".

Related to Phase 5:

We heard from a member of the ecosystem that in their early testing they observed that the reduced UA on ChromeOS would break playback: crbug.com/1298570. That was fixed in crbug.com/1302637.

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

CMA concerns

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

Stakeholder concerns

The CMA has informed Google about certain concerns expressed by stakeholders:

Topics interaction with Search ranking - The CMA has expressed concerns received from market participants regarding whether a site's opt-out of the Topics API will impact its position in Search. The Topics API is designed to support the interest-based advertising (IBA) use case, in a manner that keeps people's activity private across a free and open internet. Topics supports IBA by sharing some coarse-grained interest signals with API callers, which can be used to personalize ads. Some websites may choose to opt-out of the Topics API. The Privacy Sandbox team has not coordinated with or requested from the Search organization that they use page ranking as an incentive for websites to adopt the Topics API. Google has confirmed to the CMA that Google Search will not use a site's decision to opt-out from the Topics API as a ranking signal.

Timeline - The CMA has continued to receive feedback from stakeholders concerning the timing of the removal of third-party cookies, specifically that uncertainty over the date of third-party cookie deprecation has the potential to create a "chilling effect" on cookie-less innovations. In Q3 2022, after consultation with the CMA, Google published a revised timeline extending the window for testing the Privacy Sandbox APIs. The revised timeline is

also intended to provide greater clarity in an effort to meet requests from the ecosystem for increased transparency around the Privacy Sandbox milestones, so that they can forecast resource allocation for testing and provide feedback on the APIs. 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 with regard to their effectiveness in replacing third-party cookie functionality and meet the Development and Implementation Criteria set out in the Commitments. Additionally Google continues to signal in its messaging to partners and the ecosystem that they should invest in, and test, relevant cookie-less technologies in preparation for the deprecation of third-party cookies. In early December 2022, Google published a [blog post](#) entitled "Maximize ad relevance after third-party cookies" to educate the ecosystem about maximizing performance using a range of privacy-safe signals.

Experiments - The CMA has explained that some stakeholders have expressed interest in learning more about Google's approach towards testing and analyzing competing solutions and how these will be accounted for when conducting effectiveness testing.

Regarding competing solutions, 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.⁹

Google strongly encourages third parties, when conducting effectiveness testing of the Privacy Sandbox, to utilize their own complementary solutions and the range of privacy-safe signals available to them to optimize their own targeting and measurement functionality. This is needed to enable a realistic assessment of the likely impact of removing third-party cookies. The effectiveness testing foreseen under the Commitments is designed to evaluate the performance of the Privacy Sandbox solutions for their intended use cases at scale (in the context, as just mentioned, of other signals and technologies likely to be used) against a relevant counterfactual, which will include but is not necessarily limited to the functionality provided by third-party cookies.

Standards Development - The CMA has highlighted that some stakeholders have concerns that the pace at which the Privacy Sandbox proposals are being developed does not leave sufficient time for proper consideration by standards bodies. Google believes that standards are essential for a functioning web and is fully committed to participating in the relevant W3C processes. Multiple Chrome teams and many third parties are working in various W3C groups such as the [Private Advertising Technology Community Group](#) (PATCG), [Web Platform Incubator Community Group](#) (WICG), [Federated Identity Community Group](#), and others, to identify and work on solutions that are broadly acceptable across many browser engines.

⁹ See Google's Q2 2022 Progress Report, page 22.

The CMA has also heard some concerns with regard to the alleged influence of Google over the W3C decision-making process. Google would like to clarify that W3C community groups like WICG are open to all, contain members from across the entire industry, and have discussions that are publicly available through GitHub, without joining W3C. The internal workings of WICG are regulated by a [Charter](#). Participants choose their Chairs (and can replace them at any time) who are responsible for ensuring that the process is fair, respects the consensus of the group, and does not unreasonably favor or discriminate against any group participant or their employer. In the Q4 reporting period, important improvements to several proposals were made based on feedback from participants in the incubation process, such as [improving latency](#) and [limiting bidding to component auctions](#) in FLEDGE, along with [additional debugging reports](#) and [attribution window improvements](#) for Attribution Reporting API, as examples. Importantly, WICG is an incubation venue: decisions on web standards are not adopted in that framework, but are instead made by the Working Groups that incubations have to graduate into in order to become web standards. In those Working Groups, [any effective “votes” are taken per-company](#), and so having multiple participants does not imply multiple votes.

As the web moves away from cross-site tracking, Google needs to make sure that the new technologies it develops effectively support the needs of the ecosystem. In some cases, as standards often take time and consensus to be established, Google is testing solutions in parallel. These results will feed into the standard-setting process and allow participants to reach a more informed consensus. Recognising the differing opinions in the ecosystem, if a different mutually-agreeable standard arrives and gains consensus in the standards setting group, Google would work with the ecosystem to support a thoughtful transition to the new APIs. Google’s long-term goal remains to create interoperable standards that multiple browsers broadly support and that provide effective, privacy-enhancing solutions for targeting and measurement use cases.

Design compatibility - Concerns have also been raised that websites might break on other web browsers if they don’t also implement Privacy Sandbox APIs. We are conscious of the importance of ensuring user experience is not compromised because of cross-browser differences. However, this is an issue that is not specific to the Privacy Sandbox APIs, and it ultimately depends on how the site's developers choose to code their sites. Developers are familiar with these issues and there is already a wide range of capability support across browsers.

First-Party Sets - The CMA has received concerns that First-Party Sets bakes in unnecessary use cases permanently, in contrast to the temporary “allow lists” of other browsers. Google views the public, transparent, and easily accessible nature of the FPS list as fundamentally valuable for the ecosystem. This allows for public scrutiny of any new or existing use cases and ensures flexibility to adapt to new and changing use cases. New submissions to the canonical FPS list must be [filed as pull requests](#) on GitHub. At the same time, to avoid the risk of multiplying unnecessary use cases, we have associated domains capped at three. Concerns regarding the three domain limit have been addressed in the FPS feedback table. As we progress towards General Availability of the APIs, Google is

continuing to refine the design of its First-Party Sets proposals in response to CMA and other stakeholder feedback.

Status Meetings

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

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

Standstill

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

Compliance statement

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



COMPETITION AND MARKETS AUTHORITY
Case 50972 - Privacy Sandbox
Compliance Statement

I, Renée M. DuPree, Director, Competition Compliance of Google LLC confirm that for the three months to 31 December 2022, 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