

CMA Q4 2023 update report on implementation of the Privacy Sandbox commitments

January 2024

Contents

Summary	2
Dashboard	4
Context and framework of our assessment	5
<i>The Commitments framework</i>	5
<i>Proposed approach during the Standstill Period</i>	7
<i>Overall competition concerns</i>	8
Potential concerns and current views on the individual Privacy Sandbox tools	11
<i>Showing relevant content and ads</i>	12
<i>Measuring digital ads</i>	23
<i>Strengthening cross-site boundaries</i>	29
<i>Fighting spam and fraud on the web</i>	38
<i>Limiting covert tracking</i>	40
Other updates covering the reporting period	46
<i>Privacy Sandbox tools</i>	46
<i>Update on testing and trialling</i>	46
<i>Actions and conclusions of the Monitoring Trustee</i>	49
<i>Technical Expert</i>	50
<i>Engagement with market participants</i>	50
<i>Engagement with the ICO and international authorities</i>	54
Next steps	54
Contact details	55
Annex 1 – current proposals in the Privacy Sandbox	56

Summary

1. This report sets out the CMA's latest views on the potential impact of Google's proposed Privacy Sandbox changes (see Annex 1), based on the framework for assessment set out in the legally binding Commitments that Google made in February 2022 to address competition concerns relating to its proposals to remove third-party cookies from Chrome. It summarises stakeholder views we are aware of on the various proposals and highlights areas where competition concerns remain. We have raised these issues with Google¹ and intend to work with it to resolve our concerns between now and the point at which Google triggers the Standstill Period.²
2. Google cannot proceed with third-party cookie deprecation until our concerns are resolved. Once a resolution is achieved, Google will be able to remove third-party cookies without delay. Subject to our concerns being resolved, Google intends to deprecate third-party cookies in the second half of 2024.
3. We are setting out these views at the start of a period of testing which will provide further evidence on the likely impacts of the Privacy Sandbox tools.³ Our assessment of the Privacy Sandbox tools at the Standstill Period will combine all the evidence available to us at the time, including the results of testing, evidence provided by third parties, along with any assurances that Google provides to resolve any remaining competition concerns. This means that the views set out in this report should be understood as preliminary indications only. Nevertheless, we hope that this provides helpful context to interested parties about our initial thinking on the potential impact of the Privacy Sandbox and whether competition concerns remain.
4. Based on the available evidence, we consider that from 1 October 2023 to 31 December 2023 (the relevant reporting period), Google has complied with the Commitments. This means that in our view Google has followed the required process set out in the Commitments and is engaging with us to resolve our remaining concerns ahead of third-party cookie deprecation. However, further progress is needed by Google to resolve our competition concerns ahead of

¹ Paragraph 17.a.ii of the [Commitments](#) enables us to raise issues with Google, and for Google to work with the CMA without delay to seek to resolve concerns raised.

² Under paragraph 19 of the [Commitments](#), Google must allow for a Standstill Period of at least 60 days before third-party cookies can be removed. This period can be extended to 120 days.

³ Google's plan to disable third-party cookies for 1% of Chrome users from Q1 2024 is specifically for the purposes of facilitating testing, and it is not the start of third-party cookie deprecation which, as mentioned above, is subject to the Standstill Period and our competition concerns being resolved.

deprecation. Any developments in January 2024 will be covered in our next update report.

5. In Q1 2024, we will focus on working with Google to resolve the competition concerns we have identified in this report. We are particularly keen on resolving any remaining concerns relating to the design of the Privacy Sandbox tools and to ensure that Google does not use the tools in a way that self-preferences its own advertising services. As part of this, we are also looking to clarify the longer-term governance arrangements for the Privacy Sandbox. We would welcome comments from interested parties on our analysis of the concerns so that we can take these into account in our discussions with Google between now and the Standstill Period.
6. Feedback can be provided to us using the contact details at the end of this report by **27 February 2024**. We will update on the views gathered from external stakeholders and how discussions with Google are progressing in the report to be published at the end of April 2024. While it may not be possible for us to respond to each individual comment, raising these points means we are better able to assess the development of the Privacy Sandbox and ensure that Google meets its legal obligations.

Dashboard

Dashboard: summary of CMA view on current position, October-December 2023

Relevant section of Commitments	Compliance	Level of focus by CMA ⁴	Key actions during period	Summary of planned next steps
D - Transparency and consultation with third parties	Compliant	Higher focus	- Engagement with Google and market participants on the development of individual proposals (eg Protected Audience API) - Following up on the recently published update to our guidance on testing	- Engaging with Google and other market participants to resolve concerns relating to the development of the Privacy Sandbox tools - Following up on the recently published update to our guidance on testing
E - Involvement of the CMA in the Privacy Sandbox proposals	Compliant	Higher focus	- Encouraging testing and trialling by Google and other market participants - Engaging on design issues including approach to Related Website Sets, Protected Audience API and Attribution Reporting API	- Engaging with Google and other market participants to resolve concerns relating to the development of the Privacy Sandbox tools - Encouraging testing and trialling by Google and other market participants and engaging with market participants that intend to test
F - Standstill before the Removal of Third-Party Cookies	Compliant	Medium focus	Preparing for the standstill, including by identifying remaining competition concerns and through testing and trialling	Continuing to prepare for the standstill, including through resolving remaining competition concerns and through testing and trialling (see above)
G - Google's use of data	Compliant	Medium focus	- Solidifying overall understanding of Google's internal data control systems (particularly those relevant to paragraphs 25 and 26) - Working to ensure that necessary data use protections are fully implemented well in advance of third-party cookie deprecation	- Resolving remaining questions regarding Google's internal data control systems (with a particular focus on those relevant to paragraph 27) - Working to ensure that necessary data use protections are fully implemented well in advance of third-party cookie deprecation - Continuing to develop and finalise framework for ongoing monitoring following third-party cookie deprecation
H - Non-discrimination	Compliant	Medium focus	- Systematising recurring elements of reporting on Section H measures - Engaging with Google to understand how developments particularly around Protected Audience API and Related Website Sets align in this context - Further testing Google's internal decision-making process, particularly at key decision points - Continuing to apply technical knowledge to monitoring artifacts and logs	- Continuing to engage with Google to understand how developments particularly around Protected Audience API and Related Website Sets align in this context - Additional similar engagement regarding the API user attestation and enrolment process - Continuing to apply technical knowledge to monitoring artifacts and logs
I - Reporting and compliance	Compliant	Lower focus	Completion of regular monitoring report(s)	- Google to continue demonstrating ongoing compliance - Preparing for next monitoring report(s)

Note: this is a summary, so it cannot provide comprehensive details on all topics

Context and framework of our assessment

7. We have set out below our current views on the proposed Privacy Sandbox changes. We first summarise the framework for assessment in Google's Commitments, before highlighting the potential competition concerns which need to be resolved.

The Commitments framework

8. The Purpose of the Commitments is to address the competition concerns we identified during our Competition Act 1998 (**CA98**) investigation, namely that, without sufficient regulatory scrutiny and oversight, the Privacy Sandbox proposals could:⁵
- a. distort competition in the market for the supply of ad inventory and in the market for the supply of ad tech services, by restricting the functionality associated with user tracking for third parties while retaining this functionality for Google;
 - b. distort competition by the self-preferencing of Google's own advertising products and services and owned and operated ad inventory; and
 - c. allow Google to deny Chrome web users substantial choice in terms of whether and how their Personal Data is used for the purpose of Targeting or Measurement and delivering advertising to them.
9. The Commitments state that Google will design, implement and evaluate the Privacy Sandbox proposals by taking into account the following factors (the 'Development and Implementation Criteria'), which will inform the answer to the question of whether or not the Purpose of the Commitments, as defined above, has been achieved. The Development and Implementation Criteria are:⁶
- a. impact on privacy outcomes and compliance with data protection principles as set out in the Applicable Data Protection Legislation (**D&I A – Privacy outcomes**);
 - b. impact on competition in digital advertising and in particular the risk of distortion to competition between Google and other market participants (**D&I B – Competition in digital advertising**);

⁴ While all aspects of the [Commitments](#) are important, this column refers to the relative priorities of the CMA, and which have required a greater focus, during the course of the reporting period.

⁵ See paragraph 7 of the [Commitments](#).

⁶ See paragraph 8 of the [Commitments](#).

- c. impact on publishers (including in particular the ability of publishers to generate revenue from advertising inventory) and advertisers (including in particular the ability of advertisers to obtain cost-effective advertising) (**D&I C – Impact on publishers and advertisers**);
 - d. impact on user experience, including the relevance of advertising, transparency over how Personal Data is used for advertising purposes, and user control (**D&I D – User experience**); and
 - e. technical feasibility, complexity and cost involved in Google designing, developing and implementing the Privacy Sandbox (**D&I E – Technical feasibility for Google**).
10. Under the Commitments, Google will work with us without delay to seek to resolve concerns raised and address comments we made with a view to achieving the Purpose of the Commitments.⁷ Google will inform us of how it has responded to those comments. In practice, this means that Google will provide the CMA with assurances on the actions it has taken or will take (or refrain from) to resolve any remaining concerns.
 11. In the event that we cannot reach mutual agreement or resolve concerns within 20 working days of written notice by the CMA (unless extended by mutual consent), we may take action, including by reopening the CA98 case.⁸ We have not served any such notice to date.
 12. The Commitments also require that Google will not implement the removal of third-party cookies before the expiry of a Standstill Period of no less than 60 days after Google notifies the CMA of its intention to implement their removal.⁹ Google may increase the length of such a Standstill Period at any time between giving such notice and the period's expiry. At the CMA's request, Google will increase the length of this Standstill Period by a further 60 days to a total of 120 days.
 13. During the Standstill Period, we may notify Google that competition law concerns remain such that the Purpose of the Commitments will not be achieved.¹⁰ Google will work with us without delay to seek to resolve concerns raised and address comments made by the CMA with a view to achieving the Purpose of the Commitments. Google will inform us of how it has responded to those comments. In practice, this means that Google will provide the CMA

⁷ See paragraph 17.a.ii of the [Commitments](#).

⁸ Pursuant and subject to the provisions of section 31B(4) CA98. See paragraph 17.a.iii of the [Commitments](#).

⁹ See paragraph 19 of the [Commitments](#).

¹⁰ See paragraph 21 of the [Commitments](#).

with assurances on the actions it has taken or will take (or refrain from) to resolve any remaining concerns.

14. As part of its Commitments to the CMA, we will monitor Google's compliance with those assurances following deprecation of third-party cookies.
15. If Google fails to comply with the Commitments, including any of the assurances provided to us, the CMA may continue its investigation under section 31B(4)(b) CA98 or apply to the court for an order under section 31E CA98. In the event of a material change of circumstances, the CMA also may continue its investigation under section 31B(4)(a) CA98. Where the CMA continues an investigation under section 31B(4) CA98, the CMA's powers to impose interim measures and/or to make an infringement decision become available to the CMA again.
16. Accordingly, if Google fails to respond to our concerns or does not provide the required assurances, in principle, we could oppose the removal of third-party cookies – in which case we would expect to continue (reopen) the CA98 investigation if Google stated that it would otherwise push ahead with third-party cookie deprecation.

Proposed approach during the Standstill Period

17. Once Google triggers the Standstill Period, we expect to assess the evidence from the testing and trialling results, along with our own analysis of the potential impact of the Privacy Sandbox changes, informed by stakeholder responses.
18. In making this assessment, we recognise that the Privacy Sandbox represents a significant change for the entire ad tech ecosystem, and that the ecosystem will experience significant impacts – for example, impacts on revenue, on the cost of advertising, or on business practices due to changes in measurement and reporting. We expect the Chrome-facilitated testing period (which will run from Q1 to Q2 2024) to provide data on the direction (ie positive or negative) and potentially the scale of impacts on publishers and advertisers in particular.
19. We will consider any impacts (eg revenue loss) in the overall context of the Privacy Sandbox changes, including the potential to deliver benefits to consumers. The Commitments do not require that there be no loss of revenue to publishers and advertisers from the deprecation of third-party cookies and their replacement with the Privacy Sandbox tools. However, the Commitments require that Privacy Sandbox is implemented in a way which does not infringe competition law and minimises the impact on revenue to the extent possible,

while also considering privacy impacts and the legitimate aim of compliance with the applicable data protection legislation through reducing cross-site tracking.

20. The scale and direction of impacts on the ecosystem could change over time, as ad techs optimise their systems, retrain machine learning models using signals from Privacy Sandbox APIs and new Privacy Sandbox functionality becomes available. Our stakeholder engagement on specific challenges, like latency concerns around Protected Audience auctions on-device, suggests that some stakeholders are optimistic that they can iteratively improve over time. Our assessment will consider the scale and direction of impacts alongside any evidence on potential improvements (or degradations) that might occur.
21. Similarly, some existing ad tech business models will be disrupted where they currently rely on cross-site tracking technologies, including third-party cookies. The purpose of the Commitments is not to support specific business models. In assessing the Privacy Sandbox changes our focus will be on the likely impacts for competition and consumers overall.
22. Under the Commitments, the CMA will consult with the ICO on whether any concerns remain, including on privacy impacts (ie D&I A).¹¹ The CMA-ICO joint statement on competition and data protection explores the intersection between our regimes.¹² In the Privacy Sandbox context, it may be the case that specific examples of Google interventions to improve alignment with data protection principles have negative impacts on some ad tech firms, and advertiser and publisher outcomes. We are mindful of this risk and the need for careful consideration of these issues so that competition and data protection objectives are promoted overall to the benefit of consumers.
23. Given that work by the ICO is ongoing, we have focused below particularly on other remaining issues (**D&I B – Competition on digital advertising, D&I C – Impact on publishers and advertisers** and **D&I D – User experience**). Our overall assessment will need to consider these issues alongside any broader data protection and privacy concerns.

Overall competition concerns

24. Based on our current understanding of the APIs and concerns raised with us by stakeholders, we have identified a series of areas that could raise competition concerns. This does not mean that we currently think the Privacy Sandbox changes cannot go ahead, but it is important that the concerns are

¹¹ See paragraph 18 of the [Commitments](#).

¹² The CMA-ICO joint statement can be found [here](#) (accessed on 16 January 2024).

resolved, either through design changes, assurances from Google about action it will take or refrain from, or other evidence which resolves our concerns.

25. There are some broad, cross-cutting issues arising from or closely linked to the Privacy Sandbox proposals – for example, the fact that Privacy Sandbox tools do not support all current ad tech use cases and business practices, may be less interoperable than solutions based on third-party cookies and could create incentives for advertisers to move spend away from the open display market and into ‘walled gardens’ owned by firms with access to significant first-party data like Google, Meta, or Amazon.
26. We have identified several key concerns that Google will need to resolve ahead of third-party cookie deprecation:
 - a. Ensuring that Google does not design, develop or use the Privacy Sandbox proposals in ways that reinforce the existing market position of its advertising products and services, including Google Ad Manager (GAM). GAM is Google's integrated ad server and supply side platform (SSP), accounting for more than 90% of the display ads served in the UK.¹³ We are exploring risks around how GAM uses some of the Privacy Sandbox tools, such as the Protected Audience API, and how these concerns could be resolved.
 - b. Addressing specific design issues with the other Privacy Sandbox tools. Our detailed views on each tool are set out in the next section.
 - c. Clarifying the longer-term governance arrangements for Privacy Sandbox. In the absence of independent governance, Google currently retains significant discretion over how Privacy Sandbox works, develops over time, and the conditions for using Privacy Sandbox (eg requiring attestations). This creates self-preferencing risks.
 - d. We will also require assurances for the future development of the Privacy Sandbox tools. For example, the Commitments currently require Google to engage with industry stakeholders. We want this engagement to continue after our decision.
27. In addition, we will consider whether further restrictions may be needed on Google's use of first-party data to target and measure ads on Google's owned & operated (O&O) inventory.¹⁴ We are conscious of the risk that ad spend could move away from open display and into O&O inventory (or ‘walled

¹³ CMA, [Online platforms and digital advertising market study, final report](#), final report, July 2020, page 270.

¹⁴ The [Commitments](#) already impose some restrictions (see sections G and H).

gardens') – depending on the overall impact of the Privacy Sandbox changes which we are estimating through the current testing and trialling.

28. Finally, market participants have raised concerns that Google's aim to restrict all cross-site tracking will harm businesses seeking to provide interest-based targeting and measurement in competition with Privacy Sandbox. Although the obligations under the Commitments relate specifically to the impact of Google's introduction of the Privacy Sandbox proposals and not Google's approach towards other market participants' alternative technologies, Google's market position allows it to have a significant impact on the viability of alternative technologies that may compete with the Privacy Sandbox tools following the removal of third-party cookies.¹⁵ Both the Privacy Sandbox tools and possible third-party alternatives will need to comply with applicable data protection legislation.¹⁶
29. We have raised our competition concerns with Google and are working with them to resolve these, following the process envisaged under paragraph 17.a.ii of the Commitments. We have held a series of meetings with Google to discuss the issues in this report. We will provide an update on how Google is intending to resolve our concerns in the report we will publish at the end of April 2024.

¹⁵ See paragraph 4.324 of the [Commitments Decision](#).

¹⁶ See paragraph 4.325 of the [Commitments Decision](#).

Potential concerns and current views on the individual Privacy Sandbox tools

30. This section is organised by the function or use case that Privacy Sandbox APIs are intended to serve, and within each use case, by API. The relevant use cases are as follows:

- a. **Showing relevant content and ads:** Currently, third-party cookies and other forms of cross-site tracking allow for interest-based user profiles to be established and users to be targeted with ads corresponding to their profile (interest-based targeting). Cross-site tracking is also used to allow advertisers to retarget customers that have previously visited their website for remarketing purposes.

Google has developed two proposals (**Topics API**, **Protected Audience API**) to enable ads targeting and retargeting without third-party cross-site tracking.

- b. **Measuring digital ads:** Cross-site tracking may also be used to determine whether and how many ads have been served successfully to users (measurement), to help assess ad effectiveness by determining whether views and clicks on ads led to conversions (attribution), and to limit how often a specific user is shown an ad (frequency capping). It also supports the reporting of the outcomes of ad auctions to advertisers and publishers to facilitate payment and show performance of contracts.

Google has developed a new measurement and reporting tool (**Attribution Reporting API**) that does not rely on third-party cookies.

- c. **Strengthening cross-site boundaries:** For example, typically relying on third-party cookies, federated log-in allows the user to use a single method of authentication (eg username and password) to access different websites rather than creating a new username and password for each website, or to use one login to be signed in on many sites thereafter.

Google has developed a proposal (**Related Website Sets**) for companies to declare relationships among sites, so that browsers allow limited third-party cookies access for specific non-ads purposes such as facilitating a user-journey across several sites. Another tool (**Federated Credential Management**) allows users to log into particular sites without sharing their personal information with those

sites. A range of other boundary APIs have been developed (**Shared Storage API, CHIPS, Fenced Frames API**).

- d. **Fighting spam and fraud on the web:** Tracking a user's browsing activity across the web is a way to establish whether that user can be trusted or should be considered as conducting fraudulent or spam activities.

Google has developed a new API (**Private State Tokens**) to enable trust in a user's authenticity to be conveyed from one context to another, to help sites combat spam and fraud, without passive tracking.

- e. **Limiting covert tracking:** Other forms of web functionality, while not dependent on cross-site tracking, currently require the provision of information that is sometimes used to facilitate cross-site tracking. An example is the information provided through the user-agent string which provides information about the user's browser and device to the website that the user is visiting, and which is useful for optimising the user's viewing experience. A further example is the Internet Protocol ('IP') address, which is useful for detecting fraud and the geographical tailoring of content.

Google has developed a range of proposals aimed at limiting covert tracking without breaking currently supported use cases (**User-Agent Client Hints/User-Agent Reduction, IP Protection, DNS-over-HTTPS, Storage Partitioning, Network Partitioning, Bounce Tracking Mitigations**)

- 31. We outline the remaining concerns we have identified for each of the APIs based on the Commitments framework (**D&I B – Digital advertising, D&I C – Impact on publishers and advertisers and D&I D – User experience**).

Showing relevant content and ads

Topics API

Overview

- 32. The Topics API is intended to enable interest-based targeting.¹⁷ It uses an on-device classifier model to generate a list of topics reflecting the user's interests based on their browsing history. The topics are selected from a human-curated, publicly available taxonomy, currently containing 469 topics.¹⁸

¹⁷ An overview of the Topics proposal can be found [here](#) (accessed on 16 January 2024).

¹⁸ The taxonomy is listed on the GitHub page [here](#) (accessed on 16 January 2024).

Human curation is intended to ensure that topics are interpretable, for example ‘Arts & Entertainment’, and that sensitive topics are excluded.

33. The on-device classifier uses the site’s hostname, including subdomains. For example, the site news.bbc.co.uk is assigned ‘news’ and the site sport.bbc.co.uk is assigned three topics: ‘news’, ‘sport’ and ‘soccer’.¹⁹ The classifier only considers the hostname.
34. Every week, Chrome will calculate (locally on the user’s device) the top five topics from the user’s browsing history of sites that use the Topics API that week (epoch). When callers (including third-party ad tech or advertising providers) call the Topics API, the API will return at random for the user up to three topics in total from the top five topics for each of the last three weeks, once a topic is selected for a week, user, and top-level site, it will remain constant. Google selects ‘top’ topics, first based on their utility (‘high’ or ‘standard’), and then by their frequency count.²⁰

Potential concerns

35. Based on stakeholder feedback and our own analysis of the API, we have considered the following potential concerns under **D&I B – Digital advertising** and **D&I C – Impact on publishers and advertisers**. In the table below, we also include our provisional views on each of the concerns identified.

Potential concerns	Provisional CMA views
That the Topics API is likely to disadvantage small ad techs who have a more limited ‘reach’ and access to targeting information compared to large ad techs.	Google’s view is that the three-topic constraint applies equally to all ad techs and that ad techs are likely to supplement topics with other (eg contextual) signals. We consider that the impact will vary based on the degree to which an ad tech relies on the Topics API as a targeting signal. Some ad techs may have access to other sources of information about a user (eg through data sharing arrangements, data on logged in users or from Protected Audience interest group membership). However, unequal access to data is not a new problem, it exists today for ad techs using third-party cookies. The ‘reach’ problem is also not specific to Privacy Sandbox; ad techs with a larger reach have more opportunities to use third-party cookies. Privacy Sandbox means that ad techs have fewer options to extend their reach by sharing information with one another within the browser (eg cookie syncing, fingerprinting and bounce tracking are all limited). That does not stop them from sharing data on the server side, although this opportunity also exists today (eg via

¹⁹ Accessed via <chrome://topics-internals/> (accessed on 16 January 2024).

²⁰ See Google Developer Blog on ‘[Enhancements to the Topics API](#)’ (accessed on 16 January 2024).

Potential concerns	Provisional CMA views
	controller-to-controller data sharing agreements, data clean rooms or other means). Therefore, compared to the status quo, we do not consider that smaller players are likely to be disadvantaged with the introduction of the Topics API.
That Google will be less reliant on the Topics API than other market participants, given its access to first-party data.	Sections G and H of the Commitments already impose some restrictions on Google's use of its first-party data. The Monitoring Trustee has a continuing role to play in verifying Google's compliance with the relevant sections of the Commitments. We will consider whether additional restrictions may be needed to resolve this concern.
That Google might advantage itself by manipulating the Topics API taxonomy which it currently controls.	Google has told us it is developing robust governance arrangements for decision-making on issues relevant to the development of the APIs. Google has said that it remains interested in stakeholder feedback on the future governance of the taxonomy and discussion of how other industry bodies can play a more active role in developing and maintaining it. We consider that transitioning ownership to an external, industry-run group could resolve concerns that Google might advantage itself by manipulating the Topics API taxonomy. The timing of such transfer will need to be discussed further with Google.
That the level of granularity of the taxonomy may have an impact on the utility of the API for publishers and advertisers and on publishers' first-party data strategies.	Given the diversity of actors in the ad tech ecosystem, we anticipate that discussions on the most appropriate size and level of granularity for Topics taxonomy will continue. Striking the appropriate balance will be a key question for the future governance model.
That classification based only on hostname means that sites covering many topics contribute less useful information than niche sites. For example, YouTube is assigned 'Online Communities', 'TV & Video' and 'Arts & Entertainment'.	Google's Q3 2023 update report states that it 'previously considered offering functionality to classify sites into topics based on page content, and made the decision not to move forward based on privacy and security concerns'. We are aware of proposals from market participants that aim to balance privacy and security concerns against improving utility, for example by using permissions policies.²¹ Our current view is that classification based on hostname is a reasonable trade off, but we are open to proposals to develop the classifier model in the future.

²¹ See for example, [issue #224](#) on the Topics repository on GitHub (accessed on 16 January 2024).

Potential concerns	Provisional CMA views
Publishers are concerned that their sites could be misclassified or not assigned a topic, and want to control the topics that are associated with their sites.	Google has expressed concern about the risk of misclassification, eg where the Topics API classifier assigns a topic that the site owner considers to be incorrect.²² Our current view is that Google's response resolves the misclassification concern and agree with Google's view that allowing site owners to control classification risks incentivising site owners to game the system.
Allowing sites to selectively contribute to a user's topics could create a free-riding problem, ie that some ad techs can choose to observe topics without contributing to the set of topics stored on the user's device.	We are aware of specific stakeholder concerns relating to SSP 'free riding' and feedback on the way Google's Q3 2023 report addressed the issue. ²³ We have raised the concern with Google and will update in our next quarterly report.
That a site's decision to support (or not to support) the Topics API should not influence its Google Search ranking.	Google has confirmed to us that Google Search will not use a site's decision to opt-out of the Topics API as a ranking signal.²⁴ We consider that Google's assurance that a site's decision to support (or not to support) the Topics API will not influence its Google Search ranking should also extend to the other Privacy Sandbox tools.
Topics relies on user consent. If consent rates are low such that Topics are unavailable, there may be knock-on effects for interest-based targeting and publisher revenue.	We recognise that Google needs to request user consent for the Topics API to operate. We anticipate that the Chrome-facilitated testing period in early 2024 will provide further information about topics availability.
The one-week epoch means that topics are likely to be out of date, with implications for showing ads where the user may already have acted on their interest (eg by making a purchase).	Although reducing the epoch length could increase utility for advertisers, given the likely impact on privacy, on balance, we currently consider that an epoch of 7 days may be appropriate.

²² Google's [Q2 2023 progress report](#), p9.

²³ See [issue #92](#) on the Topics repository on GitHub (accessed on 16 January 2024).

²⁴ See page 34 of Google's [Q4 2022 progress report](#) (accessed on 16 January 2024).

36. As regards the application of **D&I D – User experience**, although there have been some positive developments in the design of user controls for the Topics API, eg users can review the topics assigned to them, concerns remain about the transparency of the information presented to users via the consent dialogue box and the extent to which users adequately comprehend and engage with the Topics choice.
37. We believe that further user research and testing of the dialogue box, using robust methodologies and a representative sample of users, is critical to resolve these concerns. Also, it is not clear if users will be prompted to revisit their choices, and the frequency of this. While repeated consent pop-ups, eg surfaced upon interacting with every publisher site, could lead to prompt fatigue and degrade user experience, we consider user research is important for informing how often these prompts should be shown to users and when. We are discussing the approach with Google.

Summary

38. Google needs to resolve our concerns for Topics API and our current view is that this could involve taking the following steps:
- a. Ensure there is adequate governance of the taxonomy. We are concerned that Google retaining governance of the taxonomy creates a risk of distorting competition between Google and other market participants. We want Google to set out a plan, with a timeline, to reassure market participants that decision-making on issues relevant to the taxonomy will be transparent and accountable to stakeholders. This could include transferring governance to an independent third party, with clear Terms of Reference to ensure that the taxonomy evolves in a way that balances utility for interest-based targeting with minimising re-identification risks.
 - b. Surface the Topics dialogue box periodically and consider approaches, based on user research, to remind or prompt users to revisit other Privacy Sandbox settings.

Protected Audience API

Overview

39. The Protected Audience API (formerly known as FLEDGE) (**PA**) is primarily intended to support remarketing and other custom audience use cases.²⁵

²⁵ For more information on PA API and how it works see Google Developer Blog, '[Protected Audience API](#)', 27 January 2022 (accessed on 16 January 2024).

Remarketing is the practice of serving targeted ads to individuals based on their activity on an advertiser's website. PA allows sites to assign users to interest groups. The browser stores information about interest groups including the name of the interest group, the group's owner and information about the interest group's configuration.

40. PA has several components, intended to work together to facilitate privacy preserving remarketing. Google has published a timeline showing the status of each component.²⁶ The timeline is high level, indicating the quarter in which Google expects the feature to be available.

Potential concerns

41. Based on stakeholder feedback and our own analysis of the API, we have considered the following potential concerns under **D&I B – Digital advertising** and **D&I C – Impact on publishers and advertisers**. In the table below, we also include our provisional views on each of the concerns identified.

Potential concerns	Provisional CMA views
Interest Group (IG) design currently excludes traffic shaping, the practice of filtering or curating bid requests to prioritise DSP responses based on some information about the bid opportunity.	Google has responded to the traffic shaping concerns by recommending that SSPs can use caching or DSPs can make increased use of Trusted Key/Value servers to address these use cases.²⁷ We recognise that traffic shaping contributes to efficient use of ad tech resources. Some ad techs are constrained by limits on the number of queries per second they can process, and traffic shaping can help them to prioritise bids. We are keen to hear further stakeholder feedback on whether Google's recommendation addresses this use case.
PA does not currently support effective IG delegation, (ie where one party assigns a user to an IG and allows another party to bid on that IG in a PA auction).	Google believes that its implementation accommodates all use cases and states that it should 'build additional support to make some use cases flow more smoothly in the future'.²⁸ We have raised the issue with Google and are seeking clarification on the timeline.

²⁶ See Google Chrome Developer Blog, '[Status of pending Protected Audience API capabilities](#)', 9 February 2023 (accessed on 16 January 2024).

²⁷ See 'Traffic shaping' in Google's [Q2 2023 feedback report](#) (accessed on 16 January 2024).

²⁸ See 'Publisher Interest Group Control' in Google's [Q3 2023 feedback report](#) (accessed on 16 January 2024).

PA auctions only allow buyers to bid on one IG. Buyers cannot combine IGs, for example to bid when a user is a member of both IG A and B.	Google has stated that PA does not support this type of ad targeting, and that combining IGs is incompatible with PA's current privacy model.²⁹ We currently agree with the approach to restricting remarketing and other custom audience use cases to one IG and will discuss our assessment of privacy issues (D&I A) in relation to the PA privacy model further in a future quarterly report.
PA auctions allow one ad auction per placement slot and each slot is treated independently. This creates challenges of competitive ad separation, ie ensuring that ads for competing brands do not appear in ad slots on the same page.	Competitor ad separation is an important industry use case. Stakeholders have suggested adding 'whole page'³⁰ or 'multi-tag'³¹ auctions to PA. Google has identified increased complexity and privacy risks associated with this feature. Our understanding is that Google is considering the feature request, and we will continue to monitor the issue. We are keen to hear further stakeholder feedback, and suggestions for other design changes that could help to address the issue.
PA auctions offer limited support for negative targeting, ie excluding some users from seeing a particular ad.	Discussions on negative targeting capabilities in PA are ongoing. Google has introduced functionality in response to stakeholder requests. However, some stakeholders continue to express concern that this does not fully address their cases. The ecosystem continues to propose improvements and Google has indicated that it will take some proposals on board.³² We will continue to monitor the issue, noting that we do not expect the Privacy Sandbox tools to replicate all the functionality currently available to ad techs using third-party cookies.
Latency for on-device auctions. Testing suggests that PA auctions may be slower, either due to device constraints (eg processing power available), auction design (eg waiting for information on the winning contextual bid before completing the PA auction) or network requests (eg fetching bidding or scoring logic).	Google proposes to address these concerns via design changes in Chrome and by publishing guidance for ad techs on optimising their approach to PA. The design changes add controls for sellers, allowing sellers to set limits on the time and resources buyers can consume.³³ The guidance includes recommended best practice for buyers and seller.³⁴ We are monitoring latency issues closely, recognising that high latency can lead to unsold ad inventory and negatively impact user experience. We expect that the Chrome-facilitated testing period will provide further data and welcome ongoing stakeholder feedback, particularly on whether the tools and recommendations Google has implemented are sufficient.

²⁹ See [issue #818](#) on the FLEDGE repository on GitHub (accessed on 16 January 2024).

³⁰ See [issue #98](#) on the FLEDGE repository in GitHub (accessed on 16 January 2024).

³¹ See [issue #846](#) on the FLEDGE repository on GitHub (accessed on 16 January 2024).

³² For example, see [issue #896](#) in the FLEDGE repository on GitHub (accessed on 16 January 2024).

³³ See 'Performance of Protected Audience Auctions' in Google's [Q3 2023 feedback report](#) (accessed on 16 January 2024).

³⁴ See Google Developer Blog, '[Improve Protected Audience API auction latency](#)'.

Moving processing to the device can raise concerns about overall page or device performance, with implications for search engine optimisation and user experience.	Our understanding is that Chrome uses separate worklets for the PA auction and page rendering. This allows page rendering to complete before the PA auction and should minimise impact on page load times. We anticipate that the Chrome-facilitated experiments period will provide further data on device performance issues. We welcome specific feedback from market participants on this issue.
GAM will not participate in PA component auctions unless it is the top-level PA seller. This means that publishers have to use GAM in order to access AdX demand.	Our understanding of the current approach is that GAM will only participate in PA component auctions where GAM also run the top-level PA auction. We have identified this as a priority area for Google to address. We are also exploring whether parties other than the publisher ad server should be able to run the top-level PA auction. Our current understanding is that some ad server functionality (eg pacing) may not be available unless the ad server runs the top-level PA auction.
PA reduces the information available to publishers compared with the status quo. Publishers will only receive information on the top-level winning bid, with no visibility over component auction winners.	PA is currently not designed to provide publishers full control over auction dynamics and data related to their advertising inventory. This design raises concerns that GAM, or any other top-level seller will receive more data and understanding of the relative value of impressions than either publishers or component sellers. We are exploring this further with Google and other market participants. Several stakeholders have proposed that the publisher's top seller should be able to share all the data (beyond price) with any component sellers the publisher may select.
Information on PA component auction winners will be visible to GAM, raising concerns about unequal access to information.	Our current understanding is that each PA component auction returns the outcome of the scoreAd function to the top-level auction. Google has informed us that information on individual component auctions never leaves the auction worklets. We are currently discussing with Google whether GAM has access to information that is not shared with the publisher.

GAM proposes to use machine learning to decide whether to trigger a PA auction. This raises concerns about a lack of transparency for publishers about how the system decides whether to trigger a PA auction and a lack of publisher control.	GAM has told us that its proposed model will optimise for total publisher revenue from all sources including direct deals, AdX programmatic auctions and revenue from other SSPs. GAM has clarified that publishers will have the option to turn off the machine learning feature when there are other sellers who want to participate in the PA auction. Our understanding is that the option is for publishers to either opt-in or out of the machine learning trigger. We are keen to hear feedback on whether this binary control addresses publisher concerns. Stakeholders have expressed concern that machine learning throttling could remove discretion from the ad tech ecosystem. Some publishers want to have the option to trigger a PA auction, and access to the information necessary to decide whether to trigger that auction. We are focusing on GAM's approach to PA as one of our top priorities and in-depth discussions are ongoing. We expect to be able to provide further updates in our next quarterly report.
Fenced Frames restricts available ad formats. They do not currently support native or video ads and stakeholders have requested native support for dynamic ad sizing.	We recognise significant stakeholder concerns around video and native ads once Fenced Frames are required. Google currently intends to require Fenced Frames no earlier than 2026. Google says that it has not yet designed a solution to render video in Fenced Frames. Stakeholders have raised support for the VAST standard as a specific concern. While Google is not obligated to support all existing standards, we are aware of the potential disruptions that a lack of VAST support could cause. Discussion on native ads and sizing is ongoing.³⁵ We are monitoring the issue and recognise the potential impact on publishers, advertisers, and users. Restricting ad formats could hinder the feasibility of dynamic content within existing native ad formats, limiting the potential for rivals and new entrants to introduce innovative advertising formats beyond walled gardens and potentially diminishing the overall user experience.

³⁵ See [issue #741](#) and [issue #311](#) on the FLEDGE repository on GitHub. See also Mutli-Ad size GitHub threads, on how the one ad-size gets decided, see [here](#) and the possibility of enabling multi-sized PA auction output, see [here](#) and implementing an additional Ad-Slot Size signal, see [here](#) (accessed on 16 January 2024).

Uncertainty about the impact of restrictions on transmitting signals from video players in Fenced Frames.	Our understanding is that video ads currently send real time signals to external systems, including for reporting purposes. The restrictions imposed under Fenced Frames will block these signals. We will continue to monitor this issue as Google implements a solution for video requirements ahead of the required use of Fenced Frames no earlier than 2026. We welcome further feedback from market participants on this use case, including on whether other Privacy Sandbox tools (eg Shared Storage) offer options to deliver the necessary signals.
PA has several components (eg Fenced Frames, Bidding and Auction Services, and so on) which will become available for testing and become mandatory on different timelines.	Google has published more detail on expected timelines, including when a feature is expected to be available for testing.³⁶ We welcome feedback from stakeholders on whether this information is sufficient to address their concerns around the lack of clarity on timing.
Google intends to deprecate event level reporting for PA auctions, no earlier than 2026. Once event-level reporting has been deprecated, the Private Aggregation API will become the only reporting mechanism available.	Google aims to prevent the use of event-level reporting for discovering the IG of individual visitors to the publisher's site, aligning with the privacy objectives of Fenced Frames. As the Private Aggregation API will become the only reporting mechanism within the PA API, the IG will be passed solely through "generateBid" and to "reportWin" functions. Our understanding is that this could reduce the information available to ad techs and could have an impact on their ability to optimise their bidding strategy. We are keen to hear further detail on the specific impacts and proposals for design changes.
PA auction design shifts data flows that were previously server to server onto the device. This raises concerns about transparency, and contractual issues (eg as ad techs have no contractual relationship with Google).	We recognise that Privacy Sandbox changes, including restrictions on access to information that is currently available, can impact ad tech business practices. We are working to identify these issues, including possible solutions.

³⁶ See Google Developer Blog, '[Status of pending Protected Audience API capabilities](#)' (accessed on 16 January 2024).

Concerns about the requirement to adopt Trusted Execution Environments (TEEs) to operate PA's server-side elements, such as the Bidding and Auctions Services and the Key/Value Server.	We discuss concerns relating to TEEs below. Google has indicated that some off-device services will be an optional extra for market participants who want to develop larger, more sophisticated models, allowing stakeholders to choose components aligning with their objectives. We are working with Google to explore the flexibility of server-to-server architecture and their impact on market participants.
URLs for loading scripts into PA auctions must have the same origin as the interest group owner.	Ad tech vendors commonly host applications on separate subdomains, moving these to the same origin could incur infrastructure costs and complicate reporting use cases. Google has indicated that design changes are possible, subject to resolving concerns around the web security model.³⁷ We will continue to monitor this issue and welcome further stakeholder feedback on prioritisation, ie is this a critical issue for the ecosystem.

42. Our concerns under **D&I D – User Experience** relate to the default enrolment into PA and the information notice for PA being shown once, immediately after the Topics API dialogue box. We are concerned that this user flow could potentially mislead users into perceiving an association between the two. Further, although users can modify their preferences for PA by navigating to the relevant settings page on Chrome, there are concerns regarding the ease and intuitiveness of this user journey. We are discussing these, and other UX concerns, with Google.
43. Although providing an active choice to opt into each API can lead to user fatigue, we believe that further user research could help to determine the optimal user flow for PA to ensure adequate user comprehension and engagement with relevant user controls.

Summary

44. Our analysis and stakeholder feedback has generated a long list of potential concerns relating to PA. Our current view is that the subset of concerns listed below are the most pressing and need to be resolved, and this could involve taking the following steps:
- Address the issues relating to GAM's approach to PA, including by ensuring that it does not distort competition in digital advertising

³⁷ See [issue #818](#) on the FLEDGE repository on GitHub (accessed on 16 January 2024).

between Google and other market participants, in a way that could reinforce its existing market position.

- b. Resolve concerns related to use of a third-party (ie other than GAM) ad server to run the top-level PA auction with GAM participating as a buyer in the PA component auction. Additionally, the need for publishers to have the same reporting visibility over component auctions as GAM.
- c. Resolve ad format concerns, specifically the use of video and native ads in Fenced Frames. Our understanding is that these formats are particularly important to ad revenue and that they are currently unsupported with Fenced Frames.
- d. Continue dealing with stakeholders' suggestions and providing support in developing solutions, where possible, regarding interest groups, negative targeting, on-device latency, and server-to-server architecture.

Measuring digital ads

Attribution Reporting API

Overview

- 45. The Attribution Reporting API (**ARA**) aims to allow ad techs to measure conversions without third-party cookies. A conversion occurs when the user takes an action (eg creating an account or making a purchase) after clicking on or viewing an ad.³⁸ Measuring conversions is necessary for several of ad tech's key functions, including budgeting, campaign reporting, optimising bidding strategies, and pricing ad inventory.
- 46. ARA supports two forms of reporting:
 - a. Event-level reporting. Event-level reports provide information about a specific ad event (like click or view). The browser stores information about ad events and conversions on-device and sends a report to the ad tech if a conversion attribution occurs. Chrome adds random delay and noise to the reports. Delay and noise are intended to protect user privacy by preventing ad techs from using event-level reports to track users across sites.³⁹

³⁸ See Google Developer Blog, '[Attribution Reporting](#)' (accessed on 16 January 2024).

³⁹ See Google Developer Blog, '[Event-level reports](#)' (accessed on 16 January 2024). ARA currently supports up to eight conversion categories for event-level reporting.

- b. Summary or aggregate reporting. Summary reports capture information about attributed conversions in a similar way as event level reports. The ad tech must first specify which ad event and/or conversion dimensions they would like to report on. When a conversion is attributed, Chrome encrypts the ad event and conversion information and sends it to the ad tech. The ad tech can batch these encrypted reports together and send to their aggregation service, a specialised server running in a Trusted Execution Environment (TEE) on the public cloud. The aggregation service aggregates the batched reports and adds privacy protections like noise. The ad tech can then retrieve summary reports from the aggregation service.⁴⁰ Chrome has proposed a minimum of 20 conversion events per aggregatable report.⁴¹

47. Chrome recommends ad techs to use summary and event level reports together, as they provide complementary information. Google Ads has published a technical explainer on how it is using ARA to measure conversions.⁴²

Potential concerns

48. Based on stakeholder feedback and our own analysis of the API, we have considered the following potential concerns under **D&I B – Digital advertising** and **D&I C – Impact on publishers and advertisers**. In the table below, we also include our provisional views on each of the concerns identified.

Potential concerns	Provisional CMA views
ARA does not support some types of attribution that are currently available with third-party cookies, for example multi-touch attribution.	It is clear that Google's ARA will not provide the same functionality as third-party cookies, given the desire to limit the amount of personal information shared. We consider that the changes Google has made to ARA following stakeholder feedback should increase its utility overall, including the move from fixed to flexible event reporting windows.
Coarser measurement may make it harder for publishers to value their ad inventory.	See above

⁴⁰ An overview of Aggregation Service for the Attribution Reporting API can be found on GitHub [here](#) (accessed on 16 January 2024).

⁴¹ An overview of contribution bounding and budgeting can be found on GitHub [here](#) (accessed on 16 January 2024).

⁴² See Google Ads Developers Blog, '[Optimally configure the Attribution Reporting API for ad measurement](#)' (accessed on 16 January 2024).

Potential concerns	Provisional CMA views
The proposed '20 event per aggregatable report' limit appears arbitrary and undermines ARA's utility.	Many of the Privacy Sandbox APIs require Google to define parameters. Our current view is that a '20 event' per aggregatable report limit is reasonable, given that attempting to measure small numbers of events via aggregate reports is unlikely to be useful. Google has published detailed guidance on tuning ARA, information on working with noise and details of how Google Ads uses ARA. ⁴³ Our current view is that this information should allow market participants to use ARA effectively.
Advertisers are currently able to adjust their spending on ad campaigns in real time. ARA imposes reporting delays and could lead to wasted spend that could otherwise have been reallocated.	We understand that ad techs can tune ARA to prioritise different types of reporting, including the move from fixed to flexible event reporting windows. We are keen to understand whether ARA can provide reporting that minimises delay for ad spend optimisation use cases. We welcome further feedback from market participants based on their experiments during the Chrome-facilitated testing period.
ARA degrades open display measurement compared with measurement capabilities on O&O ad inventory.	More sophisticated attribution may be possible on O&O inventory, for example where the ad tech has access to first party data. The Commitments impose restrictions on Google's use for first-party data (specifically Chrome browsing history and Google Analytics) for measurement on Google O&O inventory. We are considering whether further restrictions on Google's use of first-party data are necessary.
Market participants will be dependent on Google's APIs for ad measurement in future, which raises concerns about the ability to audit and verify results.	Ad verification currently relies on auditing of log-level reporting data, which will not be available under ARA. Google needs to explain how it sees ad verification use cases being addressed under the Privacy Sandbox.
Google's proposed approach to attribution differs from the approach taken by other browsers, which means that there may be limited interoperability of ARA with other solutions.	We remain concerned that lack of interoperability could harm competition by creating additional cost and complexity for businesses seeking to measure digital ads. Google needs to explain how it will continue its efforts to enhance greater interoperability of approaches to attribution and reporting over time.
Google limiting the number of different attributions per advertiser to eight conversion	We are concerned about Google limiting the number of different attributions per advertiser to eight conversion types, which may be harming advertisers who have more than eight

⁴³ See Google Ads Developer Blog on [optimally configuring ARA](#) and the Google Developer Blog on the [ARA and noise lab](#) (accessed 16 January 2024).

Potential concerns	Provisional CMA views
types, potentially harming advertisers who have more than eight types.	types. Currently, ARA limits 'trigger data' to 3 bits and this only allows eight distinct conversion types. Phase 2: Full Flexible Event-Level⁴⁴ would appear to indicate future support for up to 32 values, however there is currently no timeframe or deadline for this. The solution suggested within ARA for attributing to multiple domains is currently restricted to three domains⁴⁵ (which may be insufficient) and requires the advertiser to register these domains in the 'click' event (which would add overhead).
The lack of and need for a transaction ID where the data passes from the buy side to the sell side, enabling the two to connect.	We are aware of stakeholder requests for Google to provide a transaction ID to support attribution reporting. ⁴⁶ Our current view is to agree with Google that this could undermine the intended privacy model for ARA.
The need to seek explicit feedback from advertisers concerning modification to their commercial contracts given the changes to aggregation and attribution.	Stakeholders have expressed concerns that Google needs to seek explicit feedback from advertisers concerning modification to their commercial contracts given the changes to aggregation and attribution, specifically whether advertisers are willing to be billed on the basis of noisy or aggregate reporting. We continue to welcome feedback from market participants, particularly from advertisers on this point. We recognise that Privacy Sandbox represents a significant change for the ecosystem and reiterate that we do not expect Privacy Sandbox to deliver identical functionality to the technologies (like third-party cookies) that it intends to replace. We also recognise that Privacy Sandbox may have business impacts (eg changes to business practices).

49. As regards the application of **D&I D – User experience**, similar to PA API, our concerns stem from the default enrolment into ARA, the sequencing of the PA API and ARA information notice after the Topics dialogue box potentially leading to misperceptions about their association, and the accessibility of the ARA settings page. We believe that further user research could help resolve these concerns. We are discussing the approach with Google.

⁴⁴ An overview of Phase 2: Full Flexible Event-Level reporting can be found on GitHub [here](#) (accessed on 16 January 2024).

⁴⁵ See [issue #1048](#) on the Attribution Reporting repository on Github (accessed on 16 January 2024).

⁴⁶ See [issue #15](#) on the Attribution Reporting repository on GitHub (accessed on 16 January 2024).

Summary

50. Google needs to resolve our concerns for ARA and our current view is that this could involve taking the following steps:
- a. Secure greater interoperability and/or standardisation of approaches to attribution reporting. These currently appear fragmented with other available solutions including Mozilla/Meta's Interoperable Private Attribution and Safari's Private Click Attribution. The lack of interoperability could harm competition by creating additional cost and complexity for businesses seeking to measure digital ads.
 - b. Explain how Google sees ad verification use cases will be addressed under the Privacy Sandbox, including ARA.
 - c. Clarify the timeframe by which 32 values would be supported including the steps that will be taken to ensure that advertisers who have more than eight conversion types are not harmed.
 - d. Seek explicit feedback from advertisers on the impact of changes, particularly concerning modification to their commercial contracts given the changes to aggregation and attribution.
 - e. Clarify the governance of ARA, including the third-party coordinators, monitoring and data retention policies.

Trusted Execution Environment (TEE)

Overview

51. Google introduced TEEs to support use-cases where off-device processing is required while preserving the privacy of user-data. Google has control of the Chrome environment where it can determine the security and privacy characteristics of the browser. However, there are no such built-in controls outside the browser, which necessitates TEE-based solutions for device to server interactions that extend the functionality of Privacy Sandbox APIs.
52. TEEs are secure server configurations that are primarily secured through appropriate hardware environments (served by the cloud providers – see below). In addition, in a Privacy Sandbox context, code images and scripts are developed and maintained by Google and further secured by an attestation mechanism that ensures the TEEs have not been modified by third parties.

53. The TEEs are currently optional, however they will become mandatory at some point after third-party cookie deprecation. There are also TEEs for different contexts, such as the Key/Value and Bidding and Auction Servers for PA.

Potential concerns

54. Based on stakeholder feedback and our own analysis of the API, we have considered the following potential concerns under **D&I B – Digital advertising** and **D&I C – Impact on publishers and advertisers**. In the table below, we also include our provisional views on each of the concerns identified.

Potential concerns	Provisional CMA views
Cost and complexity of adopting TEEs. We have heard concerns from ad tech stakeholders about the significant financial and staffing resources required to adopt and maintain TEEs.	Google has published certain explainers to address this concern. ⁴⁷ We expect that it will publish further information (eg a cost explainer for K/V) to help ad techs estimate the costs associated with adopting TEEs. We recognise that some of these costs will vary, for example depending on specific deals with cloud providers that an ad tech can negotiate.
Google currently only supports two public cloud providers, Amazon Web Services (AWS) and Google Cloud Platform (GCP).	Google has committed to support other public cloud providers, based on feedback from the ecosystem. Google has stated that it will use ad tech feedback on which cloud services should be supported as a key prioritisation criterion. Although we agree this is a sensible route forward, we recommend that Google provide timelines for ensuring that cloud providers, and any others as the market develops, are supported within appropriate timeframes. At the very least, Google needs to provide objective criteria for new cloud providers. Competition in the cloud infrastructure level is subject to a separate, ongoing market investigation by the CMA. ⁴⁸
Google has limited support to public cloud, meaning that ad techs cannot run TEEs on their private cloud infrastructure.	Google has said that deploying TEEs on private cloud environments presents significant challenges. We note the strong ecosystem interest in private cloud (eg the issue was raised in Google's Q2 2023 and Q1 2023 feedback reports). We have raised the issue with Google and are seeking further clarity on the specific security challenges.

⁴⁷ See [Bidding and Auction Cost](#) explainer and [Aggregation](#) guidance (accessed on 29 January 2024).

⁴⁸ See CMA, [Cloud services market investigation](#).

Potential concerns	Provisional CMA views
Self-preferencing risk associated with running TEEs on GCP.	We are aware of stakeholder proposals for an entirely open-source solution, including an open-source TEE, that is inspectable and transparent. We will need to consider this as part of a range of solutions to mitigate against the self-preferencing risk.
Governance arrangements for coordinators.	Supporting alternative providers also requires Google to onboard coordinators for the Aggregation Service on that platform. We are concerned that delays in onboarding coordinators could have a negative impact on market participants, for example giving them less time to test their implementations and provide feedback.

Summary

55. Google needs to resolve our concerns for TEE and our current view is that this could involve taking the following steps:
- Provide deadlines and objective criteria for expanding the list of supported cloud providers.
 - Provide more public cost/benefit estimations related to deploying the TEEs for ARA and PA. This relates to both performance and investment cost. Market participants may be reluctant to invest in TEE infrastructure if it is not demonstrated that it is both performant and cost effective.
 - Seek feedback on the likely implementation and adoption costs of TEEs and potential impact to the market, particularly to smaller ad techs.
 - Provide detailed information on the governance of the third-party coordinators.

Strengthening cross-site boundaries

Related Website Sets

Overview

56. Related Website Sets (**RWS**) is a carve-out to third-party cookie deprecation intended to mitigate site breakages. RWS allows a set of domains to be declared as belonging to the same party. Google has said that RWS is not designed for ads use cases. When one site embeds another site and both are in the same RWS, Chrome will allow the embedded site to access its own cookies, which in the absence of the RWS would be blocked as being third-

party cookies, therefore tracking across the domains within a RWS will be possible. RWS consists of a 'set primary' domain and 'set member' domains.⁴⁹

57. Site owners declare related domains using one of three Google-defined subsets. The subsets are based on use cases, reflecting the purpose of the relationship between the set primary and the set member:
 - a. Country code top level domains (ccTLDs): For example, google.fr is a ccTLD for Google in France. The 'ccTLD' subset can contain an unlimited number of domains meeting the formation criteria. In practice, the number of domains is limited to 255, the current number of ICANN ccTLDs.⁵⁰
 - b. Service domains: For example, domains used to isolate sensitive functions (such as supporting authentication flow) from user-facing domains. 'Service domains' are domains that provide key infrastructure for a service. The 'service' subset can contain an unlimited number of domains meeting the formation criteria.
 - c. 'Associated' domains: Google uses the example of maintaining user journeys across distinct brand websites. RWS could enable those companies to share cross-site data between those domains, if the set formation criteria were met. RWS will automatically grant cross-site access to the first five domains listed.
58. RWS relies on the Storage Access API to facilitate cross-site access for domains in the 'associated' subset.⁵¹ The Storage Access API is subject to technical controls that Google has said will discourage the use of the 'associated' subdomain for ads use cases.
59. The list of subsets may evolve. Google has told us that examples of declarations may help Chrome and the broader web ecosystem identify additional use case patterns to possibly create new subsets or new APIs. Google lists set formation requirements by subset on the RWS GitHub repository.⁵² Google applies technical validation to RWS submissions. There is currently no validation other than the technical checks.

⁴⁹ The RWS Submission Guidelines can be found on GitHub [here](#) (accessed on 16 January 2024).

⁵⁰ The ICANNwiki 'Country code top level domain' can be found [here](#) (accessed on 16 January 2024).

⁵¹ An overview of how RWS leverages the Storage Access API can be found on GitHub [here](#) (accessed on 16 January 2024).

⁵² An overview of set formation requirements can be found on GitHub [here](#) (accessed on 16 January 2024).

Potential concerns

60. Based on stakeholder feedback and our own analysis of the API, we have considered the following potential concerns under **D&I B – Digital advertising** and **D&I C – Impact on publishers and advertisers**. In the table below, we also include our provisional views on each of the concerns identified.

Potential concerns	Provisional CMA views
Google discretion in merging RWS declarations into the canonical list.	Although only a handful of submissions have been made to date, it is clear a level of human intervention is still required in some cases. This will become a more urgent issue with scale as more submissions are made. We are concerned that this introduces the possibility of arbitrary discretion and would like Google to expand the governance framework to include a means for submitters to appeal if they disagree with a decision made by a human during the process. We have raised specific stakeholder feedback with Google.⁵³
Lack of clarity around the definition of ‘ownership’.	We have considered stakeholder feedback around issues relating to ownership and data controllership within RWS. We recognise Google’s desire to implement clear, automatable validation checks that effectively mitigate abuse. Automating checks also reduces Google’s discretion, therefore reducing the risk that Google will govern RWS in ways that risk distorting competition. Google has implemented a /.well-known/ metadata requirement that essentially defines ‘ownership’ as administrative access to the set member domains. Developers place a copy of the RWS declaration in the /.well-known/ folder on each set member domain. This demonstrates that they have access to modify files on each domain in the set, and prevents domains from being added to the set without their agreement. This removes some of the complexity in defining either corporate ownership or data controllership. We agree with Google’s approach to technical validation based on access to a site’s /.well-known/ directory.
RWS limits automatic cross-site data sharing to the first five domains in the ‘associated’ subset.	We believe that limiting auto-granted cross-site access to the first five domains in the ‘associated’ subset can reduce the risk of abuse when compared to third-party cookies. We acknowledge stakeholder feedback that it may be possible to combine data from more than five domains in a manner that complies with data protection law. Google arrived at the five-domain limit after

⁵³ For example, Movement for an Open Web (MOW)’s [blog on RWS](#) (accessed 16 January 2024).

Potential concerns	Provisional CMA views
	consultation with the ecosystem and undertaking user research and analysis. We expect Google to provide further evidence in support of this finding.
Prompting flow can be disruptive and undermine user experience.	We consider that this strikes an acceptable balance between utility and privacy. RWS is primarily aimed at preventing site-breakages post third-party cookie deprecation. Prompting for additional cross-site data sharing requests by sites and services enables user choice and intervention to prevent breakage where access to wider cross-site data sharing beyond the limits imposed by RWS is required.
Restrictions on the ability to combine data across sites disproportionately affects sites without access to logged-in users (eg news). Sites with a large proportion of logged-in users (eg Google) are less affected by the restrictions.	RWS, and Privacy Sandbox as a whole, will limit site owners' ability to share cross site data between 'associated' domains and this limitation may affect publishers' ability to build first party audience data. Some types of sites, specifically those with a high proportion of logged in users, may be less affected as they will have the option to combine data on logged in users on the server side. We are continuing to discuss the implications of this with Google. Paragraph 27 of the Commitments includes specific provisions on Google's use of Google First-Party Personal Data and Personal Data regarding user activity on sites other than those of the relevant publisher and advertiser to target and measure ads.

61. As regards the application of **D&I D – User experience**, the user must click on the 'tune' icon, then on 'cookies and site data' and then 'Manage cookies and site data' to see that the site they are visiting uses RWS. Our concerns on the RWS user flow involve whether users are sufficiently enabled to identify sites belonging to the same set, understand the reason for the set, and comprehend how their data is collected and shared across RWS member domains. We are discussing this with Google.

Summary

62. Google needs to resolve our concerns for RWS and our current view is that this could involve taking the following steps:
- Improve RWS governance, by defining and implementing clear policies relating to the Chrome team's role in defining the use case-based subsets, the set formation criteria, manually merging pull requests into the canonical RWS list and the currently undefined process for managing challenges.

- b. Address feedback on user controls.

Federated Credential Management

Overview

63. Federated Credential Management (**FedCM**) is intended to support federated identity on the web following third-party cookie deprecation, allowing users to choose which account to use to log in to a website via a dialog in the browser. Google has said that identity federation has played a central role in raising the bar for authentication on the web compared to per-site usernames and password in terms of trustworthiness, ease-of-use, and security.⁵⁴
64. Federated identity solutions currently rely on technologies such as iframes, redirects, and cookies – which provide vectors for user tracking across the web, and would be restricted by Google’s Privacy Sandbox changes. Google has proposed FedCM as a privacy-preserving solution to enable relying parties (RPs) to provide users with a choice of identity providers (IdPs) for sign-in and authentication.
65. Mozilla has recently started to prototype FedCM in its ‘Nightly’ experimental builds.⁵⁵ Apple has previously indicated its support for FedCM⁵⁶ although it has not yet been implemented in Safari.

Potential concerns

66. Based on stakeholder feedback and our own analysis of the API, we have considered the following potential concerns under **D&I B – Digital advertising** and **D&I C – Impact on publishers and advertisers**. In the table below, we also include our provisional views on each of the concerns identified.

Potential concerns	Provisional CMA views
FedCM might be implemented in a way which negatively impacts competition between IdPs, including Google’s own ‘sign in with Google’, through technical complexity of implementation.	We are working with Google and stakeholders to understand further the impact of technical complexity on IdPs which compete with Google.

⁵⁴ See Google Developer Blog, ‘[Federated Credential Management API](#)’ (accessed on 16 January 2024).

⁵⁵ An intent to prototype can be found [here](#) (accessed on 16 January 2024).

<https://groups.google.com/a/mozilla.org/g/dev-platform/c/ncmUwK1uO98/m/COhPA4ZrAAAJ>

⁵⁶ Apple’s response to a request for position on FedCM can be found [here](#) (accessed on 16 January 2024).

Potential concerns	Provisional CMA views
Google might unfairly benefit from greater use of federated ID within advertising solutions, as cross-domain signals are reduced.	Although we do not view support for a use case in which Google has an interest to be, in itself, an act of self-preferencing, we are keen to understand this dynamic further. We would be concerned if Google implemented FedCM in such a way which benefited its own IdP and will continue to monitor this risk.
FedCM might disintermediate publishers, restricting their ability to track users on their property.	We consider this risk to be low given that it is the publishers' choice whether to support federated login as opposed to managing user sign-in themselves. However, we invite views from publishers to understand if this is an outstanding concern.
FedCM might not support the broadest range of features, limiting its effectiveness.	We are encouraged by engagement on GitHub with respect to additional use cases and will continue to monitor the situation.

67. As regards the application of **D&I D – User experience**, there are concerns as to whether the user will receive sufficient transparency with respect to how their data will be used for advertising purposes.

Summary

68. Google needs to resolve our concerns for FedCM and our current view is that this could involve taking the following steps:
- Ensure that the API is not implemented in such a way that negatively impacts competition between IdPs, including Google's own service.
 - Mitigate the risk of cross-site correlation; unauthorised data usage; secondary use; RP fingerprinting; and ensure IdPs do not receive more data than is necessary for the purposes of authentication.

Shared Storage API

69. For Shared Storage API, we will discuss this API with Google in greater detail over the next quarter.

Cookies Having Independent Partitioned State

Overview

70. Cookies Having Independent Partitioned State (**CHIPS**) is intended to support the embedding of third-party services within webpages after third-party cookie

deprecation, without re-enabling cross-site tracking.⁵⁷ It enables developers to read and write cookies from cross-site contexts, such as iframes, in a strictly partitioned manner such that a cookie may only be accessed within the context of the top-level site where it was set.⁵⁸ Third parties who set partitioned cookies on separate webpages are not able to join up this information.

71. Google has said that CHIPS is necessary to support users' expectation of businesses on today's Internet and to facilitate website functionality such as:
 - a. Third-party embedded services including chat, maps, and payments;
 - b. Third-party Content Delivery Networks servicing access-controlled content which must be authorised by the first-party site; and
 - c. Embedded ads relying on per site frequency capping or user preferences.
72. Other browsers have considered measures to address these use cases. Firefox's solution involves partitioning all third-party cookies by default, while Safari previously attempted to partition based on heuristics before instead blocking all third-party cookies.
73. CHIPS takes a different approach and requires developers to explicitly opt-in, which Google has said will reduce confusion and unexpected bugs. CHIPS is being discussed in W3C's Privacy Community Group and appears to be moving towards cross-browser support, with outstanding discussion relating to performance and memory rather than security or privacy concerns.⁵⁹

Potential concerns

74. We have considered the following potential concerns under **D&I B – Digital advertising** and **D&I C – Impact on publishers and advertisers**. In the table below, we also include our provisional views on each of the concerns identified.

⁵⁷ An overview of the CHIPS proposal can be found [here](#) (accessed on 16 January 2024).

⁵⁸ Google provides the following illustrative example: 'For instance, when chatvendor.com is embedded on site A.com, it could request a "Partitioned" cookie to be set. Later, when chatvendor.com is loaded on site B.com, it cannot access the cookie and associated data set by it when it was previously loaded on A.com. chatvendor.com cannot join cookies that it sets across A.com and B.com to track users across the web, but chatvendor.com's key functionality of knowing who a user is across successive visits to a specific top-level site is still possible – without A.com or B.com having to trust chatvendor.com more than they do today'.

⁵⁹ See Chips repository on GitHub [here](#) (accessed on 16 January 2024).

Potential concerns	Provisional CMA views
That the partitioning of cookies by domain may reduce the ability of ad techs to compete on the targeting and measurement of advertising based on cross-domain tracking.	We accept that a reduction in cross-domain tracking is necessary to achieve the privacy benefits of third-party cookie deprecation. We consider this balance as part of our overall assessment of Google's proposals.
That CHIPS might be implemented in such a way which does not sufficiently enable advertising use cases.	We would like to understand further the extent to which CHIPS preserves existing advertising use cases.
That the partitioning of cookies by domain may reduce the effectiveness of tools for the targeting and measurement of advertising based on cross-domain tracking.	We accept that a reduction in cross-domain tracking is necessary to achieve the privacy benefits of third-party cookie deprecation. We consider this balance as part of our overall assessment of Google's proposals.
That CHIPS might be implemented with insufficient memory to enable the third-party services required by, in particular, small publishers.	We view that Google continues to make best efforts to ensure CHIPS is implemented with sufficient memory to support the greatest range of use cases. We are also pleased to see Google's engagement with other browsers, and willingness to make changes in order to progress CHIPS towards standardisation.
That CHIPS may impact the ability of publishers to offer SSO sign-in services based on authenticated embeds.	We note Google's intent to support authenticated embed use cases through Storage Access API with user prompts and continue to monitor. We would also like to understand the extent to which FedCM mitigates this concern by enabling SSO sign-in use cases.

75. Currently, we do not have any outstanding concerns in relation to the application of **D&I D – User experience**.

Summary

76. Google needs to resolve our concerns for CHIPS and our current view is that this could involve taking the following steps:
- Demonstrate that CHIPS preserves the effectiveness of legitimate advertising use cases such as frequency capping and ads personalisation.
 - Ensure that CHIPS is implemented with sufficient memory to enable third-party applications relied upon by publishers.

- c. Ensure that the ability of publishers to offer single sign-in services via authenticated embeds is preserved.

Fenced Frames

Overview

77. Fenced Frames aims to enforce a boundary between a webpage and any cross-site content it embeds, such that user data is not able to be joined up between the two sites. Under Google's PA proposal, Chrome renders the winning ad in a Fenced Frame. The requirement to render winning ads within Fenced Frames will be enforced no sooner than 2026.⁶⁰
78. Google is continuing to make gradual progress in enabling various Fenced Frames solutions, illustrated by the increased GitHub Explainer updates from October 2023. Fenced Frames does not support the same use cases as iframes currently. For example, PA supports video rendering using a mechanism that relies on iframes, and Google has not yet designed a solution that is compatible with Fenced Frames, which could significantly impact advertisers' revenue.

Potential concerns

79. Based on stakeholder feedback and our own analysis of the API, we have considered the following potential concerns under **D&I B – Digital advertising** and **D&I C – Impact on publishers and advertisers**. In the table below, we also include our provisional views on each of the concerns identified.

Potential concerns	Provisional CMA views
Fenced Frames does not currently support use cases such as native and video advertising, which may impact the ability of publishers to effectively monetise their content.	Google should implement changes to enable these key use cases before requiring ads to render in Fenced Frames.

80. Currently, we do not have any outstanding concerns in relation to the application of **D&I D – User experience**.

⁶⁰ The timelines of pending PA API capabilities can be found [here](#) (accessed on 16 January 2024).

Summary

81. Google needs to resolve our concerns for Fenced Frames and our current view is that this could involve taking the following steps:
- Explain how the API should work with other tools like Shared Storage API and PA (ie will Fenced Frames be an effective control?).
 - Explain how it proposes to address limitations of Fenced Frames, which currently does not support key use cases (native and video ads).

Fighting spam and fraud on the web

Private State Tokens

Overview

82. Private State Tokens (**PST**) enables trust signals to be transmitted between websites to determine whether a user is trustworthy or engaged in spam or fraud without allowing the user's identity to be discovered across sites. Instead, the PST aims to enable sites to collaborate in segmenting users into 'trusted' and 'untrusted' categories. To do so, a website that has already established a user's trustworthiness would be able to issue that user's browser with trust tokens.⁶¹ These tokens could then be redeemed on other websites establishing trust without identifying the user or providing information on the origin of the token. The tokens themselves will allow for limited information to be communicated.

Potential concerns

83. Based on stakeholder feedback and our own analysis of the API, we have considered the following potential concerns under **D&I B – Digital advertising** and **D&I C – Impact on publishers and advertisers**. In the table below, we also include our provisional views on each of the concerns identified.

Potential concerns	Provisional CMA views
PST could centralise Google's power by requiring sites to rely on Google to determine whether a user should be trusted.	While any entity can become a PST issuer, we believe it is conceivable that the main issuers will be well-known sites that most people visit. Given that Google owns several domains that are among the most visited sites, it is in a strong position

⁶¹ This website is known as the 'issuer'. Any website can issue trust tokens.

Potential concerns	Provisional CMA views
	to become a prominent and trusted issuer that is relied on by many sites.
Google could abuse its position as a dominant PST issuer.	To mitigate the risk to competition of Google becoming a dominant issuer of PST tokens, we recommend that Google provide policy or technical safeguards that would prevent it from abusing its position. This could be enforced through the registration and governance mechanisms that have yet to be clarified in the PST proposal. We would particularly welcome governance policies that specify why certain issuers might be disallowed from issuing PST tokens.
There will not be enough choice of PST issuers.	Our understanding is that Google has already provided demos and guides to help with setting up and running an issuer, but this does not guarantee that there will be enough competition and enough choice of issuers that are broadly trusted. To obtain greater assurance on this, Google could provide a target for how many PST issuers are expected to exist (and be actively used to redeem tokens) by the time third-party cookies are deprecated.

84. As regards the application of **D&I D – User experience**, our understanding is that users can opt-out or otherwise negate the functioning of PSTs using an ‘Auto-verify’ feature in Chrome. Given that PST currently permits a non-exhaustive range of use cases, we are concerned that stakeholders may use PST data for purposes besides verification. We consider this to be problematic, as users will not be made aware of the other purposes for which their data may be used.

Summary

85. Google needs to resolve our concerns for PST and our current view is that this could involve taking the following steps:
- Provide clear registration and governance criteria for PST issuers.
 - Provide assurance that there will be enough competition and enough choice of PST token issuers that are broadly trusted.
 - Clarity as to whether PST tokens can be used for purposes besides verification. If so, we request Google to clarify how users will be informed of the other ways in which their data may be used.

Limiting covert tracking

Bounce Tracking Mitigations

Overview

86. Bounce Tracking Mitigations (**BTM**) is intended to address cases where sites use a ‘stateful bounce’ to identify users across different sites. A ‘stateful bounce’ allows sites to replicate the cross-site tracking functionality of third-party cookies. For example, the user navigates to Site A, Site A redirects the user to Site B, Site B accesses state (eg sets a cookie, accesses local storage, and so on) and redirects the user again either back to Site A or to another site.
87. These redirects can happen quickly, and users may not be aware of them. Google’s implementation of BTM relies on user interaction. If the user has interacted with the site that they are redirected to (Site B in our example above) within the last 45 days, the ‘stateful bounce’ will be allowed, otherwise the state (eg the cookie set by Site B in our example) will be deleted.
88. Google has identified some use cases that rely on stateful bounces that will continue to work because they involve user interaction. These use cases include: (1) federated authentication, (2) single sign on and (3) payments.⁶² Google has invited specific feedback on whether user interaction is the most appropriate signal to indicate that the stateful bounce is part of a use case that should be supported under BTM.⁶³

Potential concerns

89. Based on stakeholder feedback and our own analysis of the API, we have considered the following potential concerns under **D&I B – Digital advertising** and **D&I C – Impact on publishers and advertisers**. In the table below, we also include our provisional views on each of the concerns identified.

⁶² An overview of out-of-scope use cases can be found [here](#) (accessed on 16 January 2024).

⁶³ See [issue #24](#) on the Navigation-based Tracking Mitigations repository on GitHub (accessed on 16 January 2024).

Potential concerns	Provisional CMA views
There is a risk that the current implementation of BTM will disadvantage competitors that rely on legitimate use of browser storage.	Our understanding is that the user interaction requirements of BTM may undermine the ability of competitors and other stakeholders in the Privacy Enhancing Technology (PET) market to use redirect flows for legitimate purposes. Google's response to this issue ⁶⁴ recommends using additional consent flows or the Storage Access API. However, we were informed by stakeholders that these options would add an unacceptable level of user friction that would break their use case. We are continuing to discuss with Google how these concerns may be resolved. This includes exploring an alternative implementation of BTM that could use list-based approaches to identify trackers, or adapting the Shared Storage API in a way that would allow legitimate use of browser storage by PETs.
Sites that are regularly visited by users (ie more than once every 45 days) will still be able to use bounce tracking.	We are concerned that Google would be able to circumvent the protections provided by BTM because of the large volume of user interactions that Google sites receive. This could give Google an advantage over competitors that have smaller audiences. We have asked Google to confirm that it will not use bounce tracking outside of their accepted use cases (eg login and payments).
There is a risk that BTM will reduce competitors' ability to use link decoration.	We have received concerns that Google's implementation of BTM tampers with URL strings. Our understanding is that link decoration is not affected in the current implementation of BTM or other Privacy Sandbox proposals.
There was insufficient industry consultation before the release of BTM.	Industry stakeholders have had the opportunity to comment on Google's BTM proposal since it was announced publicly in September 2022. Possible avenues for stakeholder engagement include the corresponding GitHub repository ⁶⁵ and relevant W3C groups.

90. As regards the application of **D&I D – User experience**, we consider that Google has taken fair precautions to ensure that BTM does not adversely affect user experience.

⁶⁴ See [issue #64](#) on the Navigation-based Tracking Mitigations repository on GitHub (accessed on 16 January 2024).

⁶⁵ See the Navigation-based Tracking repository on GitHub [here](#) (accessed on 16 January 2024).

91. We will continue to monitor how BTM impacts user experience in technologies that currently rely on redirection and browser storage for legitimate use cases (eg PETs and authentication).
92. While BTM will not be enforced until after third-party cookie deprecation, BTM is currently available for testing/use by anyone who has already blocked third-party cookies. This gives stakeholders and users an opportunity to monitor and feedback the impact of BTM on user experience before it becomes fully operational.

Summary

93. Google needs to resolve our concerns for BTM and our current view is that this could involve taking the following steps:
 - a. Resolve concerns about BTM's current user interaction requirement.
 - b. Confirm that the current BTM implementation would not allow sites with a large first party presence (like Google) to use bounce tracking. We are waiting for Google's confirmation on this point.

User-Agent Client Hints/User-Agent Reduction, IP Protection, DNS-over-HTTPS, Storage Partitioning and Network Partitioning

Overview

94. The purpose of **User-Agent Client Hints (UA-CH)**, which follows from **User-Agent Reduction (UAR)**, is to limit passive fingerprinting of users, limiting the amount of information the browser automatically delivers about the user to the web server it interacts with through the User-Agent String. The User-Agent String is transmitted as a request header in every HTTP exchange between client and server. The process is generally opaque to users. UA-CH therefore enforces a model whereby the server must actively request details about the client (eg device model) rather than passively receive them.
95. IP Protection is a proposed privacy feature in Chrome that aims to avoid sharing a user's real IP address with third parties. Under the current proposal, a privacy proxy will be used to anonymise eligible users' IP addresses.⁶⁶ Google will use two proxies where the first is run by Google and the second by an external content delivery network (CDN). Google's aim is to (i) stop a destination origin from seeing a user's original IP address and (ii) prevent the proxy and network intermediary from seeing traffic content.

⁶⁶ An overview of the IP Protection proposal can be found [here](#) (accessed on 16 January 2024).

96. DNS-over-HTTPS is a protocol that encrypts Domain Name System (DNS) queries and responses by encoding them within HTTPS messages. This helps prevent attackers from observing what sites users visit or sending them to phishing websites.
97. Storage Partitioning will isolate some web platform APIs used for storage or communication if used by an embedded service on the site, ie in the third-party context.
98. A browser's network resources, such as connections, DNS cache, and alternative service data are generally shared globally. Network State Partitioning will partition much of this state to prevent these resources from being shared across first-party contexts. To do this, each request will have an additional 'network partition key' that must match in order for resources to be reused.

Potential concerns

99. Although we do not currently have concerns about UA-CH/UAR, we are keen to ensure that Google does not remove or limit access to critical hints in the future.
100. As regards IP Protection, based on stakeholder feedback and our own analysis of the API, we have considered the following potential concerns under **D&I B – Digital** advertising and **D&I C – Impact on publishers and advertisers**. In the table below, we also include our provisional views on each of the concerns identified.

Potential concerns	Provisional CMA views
Google may continue to benefit from user activity data while limiting competitors' access to the same data.	We may suggest that Google removes itself from the first 'hop' and use a second independent third-party instead. We will also need to understand whether data collected through sign-ins can be used in Google advertising and whether requiring user sign-in for IP Protection could be replaced by alternative means for user authentication.
Google's ability to control the inclusion of ad tech rivals on this list could advantage its ad tech services, especially if they are not subject to the same restrictions in the future.	Google will need to provide further detail on the governance process. We may suggest that independent third-party governance be required to ensure fairness and transparency.

Potential concerns	Provisional CMA views
Competition between providers of VPN services may be foreclosed.	This is beyond the scope of the Commitments but an issue that we will consider where appropriate.
Publishers and advertisers that rely on IP addresses for geographically targeting and personalising content will be forced to offer a worse service.	The provisions for GeoIP within IP Protection proposals will allow ad tech and publishers to continue to optimise content to approximate geographic location. As with the Topics API, there may be no 'right' level of granularity for all market participants. We will need to consider loss of precision in targeting against privacy benefits.
Publishers and advertisers may be less able to effectively identify fraudulent activity.	We will need further specifications from Google on how block lists or other options will be applied.
Google may provide insufficient notice for ad techs to implement alternative solutions with their publishers and test and comment back on proposals.	We will need Google to clarify and seek feedback on sufficient notice to be provided to ad techs to implement alternative solutions with their publishers and test and comment back on proposals.

101. For **DNS-over-HTTPS**, **Storage Partitioning** and **Network Partitioning**, we will discuss with Google in greater detail over the next quarter.

Summary

102. Google needs to resolve our concerns for UA-CH/UAR and IP Protection and our current view is that this could involve taking the following steps:
- a. For **UA-CH/UAR**, provide assurances that it will not further remove or limit access to critical hints in the future.
 - b. As regards **IP Protection**:
 - i. Resolve concerns that Google may obtain information from IP Protection that it can use for advertising purposes, especially in relation to Google's control of the first 'hop' and user sign-in.
 - ii. Ensure transparency of user controls including the ability for users to revisit their choices about whether to opt-in to, or opt-out of, IP Protection.
 - iii. Ensure an adequate governance process to ensure fairness and transparency.

- iv. Provide specifications on how block lists or other options will be applied.
- v. Need for sufficient notice to be provided to ad-techs to implement alternative solutions with their publishers and test and comment back on proposals.

Other updates covering the reporting period

Privacy Sandbox tools

IP Protection

103. In Q4 2023, Google began Phase 0 of its plans for testing and implementing IP Protection.⁶⁷ This phase involves initial functional testing of the 1-hop proxy on Google owned domains only. Some stakeholders raised concerns that Google was prematurely implementing IP Protection. In response, we asked Google to publish an update clarifying the various phases in the development and implementation of the proposal. Google has now made this available in its Q4 2023 report. The Commitments do not require Google to test and evaluate the impact of IP Protection at this stage.

Privacy Budget

104. Google has announced that Privacy Budget is no longer being actively considered as part of its Privacy Sandbox proposals. As set out in the Commitments Decision, Privacy Budget was among those proposals aimed at combating fingerprinting.⁶⁸ Google planned to use a browser-assigned information budget to limit the data provided to individual websites. We will continue to monitor any further plans that Google puts into place for addressing fingerprinting across the web.

Web Environment Integrity

105. Over the last quarter, Google has also decided not to move ahead with its proposal for Web Environment Integrity (**WEI**).⁶⁹ This was designed to help publishers detect invalid traffic by evaluating the authenticity of the user's device, software stack and behaviour. Due to concerns raised over its potential risk to competition, we were considering how WEI could be addressed under the scope of the Commitments. However, given that the proposal has been abandoned, we will not be taking any further action.

Update on testing and trialling

106. During the reporting period, we continued to work through technical aspects of Google's internal testing in readiness for the launch of the testing period, **during Q1 and Q2 2024**. Google has previously run several internal tests of

⁶⁷ An overview of Google's plans for implementing IP Protection can be found [here](#) (accessed on 16 January 2024).

⁶⁸ See paragraph 3.23 of the [Commitments Decision](#).

⁶⁹ Google has updated its original explainer [here](#) (accessed on 16 January 2024).

the Privacy Sandbox targeting and measurement APIs in isolation (ie Topics, FLEDGE – now PA, and ARA).⁷⁰ These tests were primarily aimed at understanding the functionality of the APIs and any early signs of their effectiveness for Google.

107. As we move into the testing period over Q1 and Q2 2024 (see Figure 1), Google will launch a combined end-to-end experiment utilising 1% of traffic for which Chrome will disable third-party cookies as part of its Mode B testing initiative. Using Mode B traffic will ensure that third-party cookies are not used in auctions intended for testing the Privacy Sandbox tools. This test will seek to estimate the potential direction and scale of impacts on Google and the advertisers and publishers who rely on its advertising services.
108. We understand it is not possible to entirely replicate what the digital advertising market might look like following deprecation of third-party cookies in Chrome, not least because market participants require time to fully develop their systems for using the Privacy Sandbox tools. (We describe third party feedback on this below). As such, we intend to view results from any quantitative tests alongside a wider qualitative and technical assessment of the Privacy Sandbox tools.
109. Given that Google does not represent the entire digital advertising market, this wider evidence base will include testing results from third-party market participants. We encourage market participants to conduct tests during the testing period in line with our guidance and submit their results to us in advance of our assessment, ideally as soon as practically possible.⁷¹
110. We have engaged with a wide variety of market participants who have begun testing or intend to run tests of the Privacy Sandbox tools during Q1 and Q2 2024. We have also reached out to participants across the ecosystem to understand any barriers to engagement more broadly.
111. Although a number of market participants intend to run experiments and submit results to the CMA, our engagement has highlighted some concerns market participants have around the comprehensiveness of any tests, either from Google or third parties. These concerns centre around two main themes:
 - a. The readiness of the ecosystem to test the Privacy Sandbox in a way that might replicate the competitive environment with third-party cookies deprecated in Chrome. For example, some market participants have told us that a lack of engagement in testing the Privacy Sandbox

⁷⁰ See for example paragraphs 17-19 of our [Q2 2023 Update Report](#).


⁷¹ See our [guidance to third parties on testing](#) and corresponding [additional guidance](#) on specific practical elements of testing.

tools might artificially lower demand and supply for ad impressions in experiments.

- b. Small sample sizes and the ability to identify impacts of the Privacy Sandbox. As part of Mode B, Google has disabled TPCs on 1% of traffic in Chrome to form the treatment group. Some market participants have suggested this might not provide sufficient sample size to capture precise impacts the Privacy Sandbox tools might have on their business.

112. We understand the implications these concerns for experiments and intend to view any results in their context. In addition, and as discussed mentioned above, in our assessment we will not view testing results in isolation but alongside a wider evidence base. This will include evidence gathered from across the digital advertising ecosystem on how the Privacy Sandbox tools might affect their business. We are also asking third party testers to submit a range of quantitative and qualitative information with their results and will continue to engage with market participants on their testing plans in order to understand their tests and results.
113. We recognise that it will continue to take market participants time to fully integrate the Privacy Sandbox tools with their systems. Figure 1 below shows the testing period will run until the end of Q2 2024 (during which third party testers can submit results to the CMA). **However, it would be beneficial to our assessment if market participants who are able to test early during this period do so and submit their results in advance of this deadline.**

Figure 1: A visualisation of the testing timeline



114. For those ready, results can be submitted to privacysandbox@cma.gov.uk . Over the next period we will continue to discuss these plans with market

participants. We encourage market participants to also contact us about their plans to test the Privacy Sandbox as soon as possible. We will also continue to engage with Google on the practical and technical elements of their combined end-to-end experiment is launched so it can provide informative results.

Actions and conclusions of the Monitoring Trustee

115. The Monitoring Trustee has not informed the CMA of any instances of Google being non-compliant with its obligations under the relevant paragraphs of the Commitments.
116. Although the Monitoring Trustee's quarterly report represents a snapshot in time, Google is subject to continuous monitoring for the duration of the Commitments. Therefore, monitoring activities may be reported on as in progress or otherwise in the process of discussion, negotiation, investigation, or consideration, with a future road map of monitoring work at any given time.
117. During the reporting period, the Monitoring Trustee has overseen Google's activities relating to paragraphs 25-27, 30-31, and 33 of the Commitments. These activities are largely a continuation of, and build upon, the work undertaken in the previous periods, including:
 - a. Continuing to review compliance artifacts around internal decision-making processes (eg logs and records) to test whether Google's internal processes are being followed in practice.
 - b. Building a deeper understanding of Google's internal data control systems in order to robustly test Google's proposals to address its commitments on Chrome browsing history, Google Analytics data, and ad inventory on websites not owned and operated by Google. These commitments only apply after Chrome ends support for third-party cookies, but we are working to ensure that these controls are fully implemented well in advance of third-party cookie deprecation.
 - c. Developing plans to investigate data flows within Google to ensure that the data controls are effective in practice (eg addressing potential risks arising from data use from any secondary storage locations).
 - d. Reviewing Google's proposals for the new technologies and the risk that these could self-preference Google through their design, development or implementation. This has included scrutinising Google's Key Design Decisions to test their compliance with Section H of the Commitments.

118. As explained below, the Monitoring Trustee has been working closely with the Technical Expert, as well as with the CMA. Submissions (or extracts of submissions) from stakeholders which are relevant to multiple elements of the compliance regime are frequently shared between the CMA, Monitoring Trustee, and Technical Expert to ensure that they are fully addressed.

Technical Expert

119. As mentioned in previous update reports, the Technical Expert aims to support the Monitoring Trustee by providing the following skills which are vital for effective monitoring of the Commitments:
- a. Analysing Google's data access and flows;
 - b. Analysing technical access controls and security; and
 - c. Providing general ad tech expertise and advice.
120. We have also continued our direct dialogue with the Technical Expert. Discussions have focused primarily on market trends and issues concerning the design and implementation of Google's Privacy Sandbox proposals. We have taken account of views and comments from the Technical Expert in our ongoing discussions with Google on the design and proposed implementation of the Privacy Sandbox tools and in identifying the remaining concerns described in the section above.

Engagement with market participants

121. We are continuing to engage with market participants in the wider online advertising ecosystem to ensure that we become aware of, and understand, concerns about the Privacy Sandbox tools and their impact.
122. Our own stakeholder engagement is not intended as a substitute for market participants' direct interactions with Google, and we would encourage participants to raise substantive concerns through existing channels including W3C. Google is required under the Commitments to respond to reasonable views and suggestions, as summarised in Google's quarterly report which is published alongside this document. It is important that Google responds substantively to feedback, and we will highlight to Google where we do not consider that it has provided an adequate response and ensure that it does so.
123. Since the publication of the CMA's last report, in Q4 2023, our engagement has had a particular focus on encouraging and guiding industry testing,

following the publication in October 2023 of an additional guidance note⁷² to market participants considering testing. We have also sought to identify and understand outstanding stakeholder concerns related to the design of Google's proposals, in addition to monitoring developments towards standardisation in W3C. Concerns raised throughout the stakeholder engagement process have been raised with Google, and directly informed our role overseeing the design and implementation of its proposals.

124. Details of the concerns raised by market participants related to the specific APIs have been included in the section above. Other concerns raised have included the following:
- a. **Technical specifications for Privacy Sandbox lack sufficient detail and can be inconsistent with other Google developer communications and blog posts.** We agree it is important that technical specifications are clear and consistent, and have raised this concern with Google.
 - b. **Google might unfairly benefit from access to data from sources such as 'Google Analytics' and the 'x-client-data HTTP request header' after the removal of third-party cookies. Additionally, Google Ads Data Manager, which builds upon Customer Match and Enhanced Conversions, might place Google in a preferential position after the removal of third-party cookies.** Google has committed that, after third-party cookie deprecation, it will not use a user's personal data from a customer's Google Analytics account in its ads systems to track that user for the targeting or measurement of digital advertising.⁷³ We are considering the extent to which further restrictions on first party data sharing might be required to avoid Google gaining an unfair competitive advantage.
 - c. **Chrome's 'Limit Covert Tracking' proposals might prevent alternative PETs from performing privacy-preserving cross-domain tracking.** We believe it is important for Google to provide greater certainty over the shape of these proposals, to enable businesses to plan and seek investment without fear of future foreclosure. We have provided further, specific thoughts in relation to BTM in the relevant section above.
 - d. **Google previously committed to not imposing a penalty in search rankings on sites which opt out of Topics API. This should be**

⁷² See CMA's [Additional guidance on industry testing](#), October 2023.

⁷³ See paragraph 26 of the [Commitments](#).

extended to apply to all Privacy Sandbox technologies. We have raised this with Google.

- e. **There are insufficient venues for public discussion of the Privacy Sandbox APIs outside of W3C.** The Commitments require Google to take into consideration reasonable views ‘including (but not limited to) those expressed in the W3C’. We believe W3C remains a key venue for Privacy Sandbox industry discussions and encourage ongoing standardisation efforts. However, we are also clear that it should not be the only means through which Google considers feedback on the proposals. Further details on Google’s approach to gathering stakeholder feedback, including a link through which to submit feedback, can be found on its dedicated [webpage](#).
- f. **While Google is limiting changes made on experimental traffic with limited exceptions during the experiment to avoid interruptions, there are some concerns that it appears to be continuing to make changes at the request of Google Ads.** We are clear that Chrome providing preferential access to Google Ads would constitute a breach of the Commitments. We have not seen direct evidence to support this claim but have raised it with Google. More widely, although we have agreed with Google that it should not make any major design changes throughout the ongoing testing period, this does not prevent them discussing potential future changes with industry.
- g. **Google should publish information on the latency impacts of Privacy Sandbox changes and outline acceptable performance service levels as part of its technical specifications.** We are conscious of concerns raised by industry that certain elements of Privacy Sandbox, for example on-device PA auctions, may not deliver acceptable latency under real world conditions. We expect to receive quantitative data on latency impact from testing currently being carried out by Google and industry and will take this into account as part of our overall assessment of the effectiveness of the proposals.
- h. **Smaller publishers and advertisers may be less able to mitigate any revenue loss from third party cookie deprecation and the potential corresponding shift from open display towards direct deals.** We are aware this potential effect and will consider it as part of our overall assessment of the competition impact of the proposals.
- i. **The Privacy Sandbox changes undermine ad techs’ ability to deliver brand safety use cases. For example, limiting**

communication between the page context and PA auctions can make it more difficult to ensure that ads do not appear alongside certain content types. We recognise that Privacy Sandbox will require some changes to ad tech operating models. We are working with stakeholders to identify critical use cases, understand whether and how they can be supported and take that information into account in our overall assessment of the competition impact of the proposals.

- j. **The enrolment and attestation requirements attached to use of the Privacy Sandbox APIs may limit the ability of ad techs to integrate Privacy Sandbox within their identity solutions.** We understand that Google's enrolment and attestation requirements play an important role in delivering the privacy provisions of its Privacy Sandbox proposals. We would, however, be concerned if these were implemented in such a way which benefited Google or placed unnecessary restrictions on the ability of ad techs to conduct business and compete.

125. In addition, several stakeholders have alleged specific breaches of Google's Commitments:

- a. **Google launched features of Privacy Sandbox in the second half of 2023 – namely, BTM and Storage Partitioning – without providing sufficient industry consultation or assessment.** It is our view that in these cases Google gave industry sufficient notice of its intentions, including via online developer explainers. Additionally, any relevant obligations associated with the Standstill Period crystallise just before Google notifies the CMA of its intention to implement removal of third-party cookies, not when BTM or Storage Partitioning are launched. Consistently with this, the impact of both these technologies will be considered as part of our overall assessment of the proposals during the Standstill Period.
- b. **Google prematurely began rolling out IP Protection without adequate notice or testing.** We asked Google to respond to this point and heard that its phase 0 testing plans involved only initial functional testing of the 1-hop proxy on Google owned domains. It is our view that the Commitments do not require Google to test and evaluate the impact of IP Protection at this stage. We have asked Google to publish an update clarifying the various phases in the development and implementation of the proposal. Google has now made this available in its Q4 2023 report.
- c. **Google is not substantively addressing stakeholder feedback.** Google's Commitments require it to 'take into consideration reasonable

views and suggestions' and provide a summary of how it has addressed this feedback on a dedicated microsite and within its quarterly reports to the CMA. We have not seen evidence to date that would suggest Google has routinely failed to comply with the reporting element of its obligations. As described in this report, we are continuing to discuss with Google how it can make further changes in response to stakeholder feedback to resolve our concerns.

126. Our focus for stakeholder engagement over the next quarter will be on guiding ongoing industry testing of the Privacy Sandbox APIs and considering views received in response to the concerns outlined in this report.
127. Given the global nature of Google's developments, we welcome feedback from organisations both within and outside the UK.

Engagement with the ICO and international authorities

128. We have continued to work together closely with the ICO in implementing the Commitments. The ICO's role has included:
 - a. Participating in discussions with us and Google on the development of the Privacy Sandbox tools, analysing data protection impacts with a specific emphasis on user controls and assessing compliance with data protection legislation;
 - b. Continuing to work with us on plans for the wider assessment of the Privacy Sandbox tools, including assessing privacy impacts; and
 - c. Engaging with market participants on proposed alternative technologies to third-party cookies and similar advertising technologies.
129. We have also continued to engage with our international counterparts and data protection authorities on the implementation of the Commitments in an effort to identify any issues of common concern and ensure consistency of approach.

Next steps

130. Over the next three months, we will focus on working with Google to resolve the concerns we have identified. Now is the time for Google to focus on making changes to ensure our concerns are resolved ahead of the Standstill Period.
131. We are planning to publish our next update report and Google's quarterly update in April 2024.

Contact details

132. We would welcome views from interested parties on this report, as well as on any other relevant publications (eg Google's own quarterly reports). The relevant contact details are:

- a. **CMA:** privacysandbox@cma.gov.uk; matthew.allsop@cma.gov.uk; angela.nissyrios@cma.gov.uk; and chris.jenkins@cma.gov.uk.
- b. **Monitoring Trustee (including communications for the Technical Expert):** trustee.services@ing.com; matthew.hancox@ing.com; and david.verroken@ing.com.
- c. **Google:** [Feedback - Chrome Developers](#).

Annex 1 – current proposals in the Privacy Sandbox

At the time of publication, the list of proposals in the Privacy Sandbox include:

- 1. Use Case: Fight spam and fraud on the web**
 - (a) Private State Tokens
- 2. Use Case: Show relevant content and ads**
 - (a) Topics
 - (b) Protected Audience
- 3. Use Case: Measure digital ads**
 - (a) Attribution Reporting
- 4. Use Case: Strengthen cross-site privacy boundaries**
 - (a) Related Website Sets
 - (b) Shared Storage
 - (c) CHIPS
 - (d) Fenced Frames
 - (e) Federated Credential Management
- 5. Use Case: Prevent covert tracking**
 - (a) User Agent Reduction (including User-Agent Client Hints)
 - (b) DNS-over-HTTPS
 - (c) Storage Partitioning
 - (d) Network State Partitioning
 - (e) IP Protection (previously Gnatcatcher)
 - (f) Bounce Tracking Mitigations