Response to Ofcom: Fraudulent Advertising Codes Consultation

Volume 4, Section 5 (Ad libraries): public libraries, enhanced access and related points

Submitted by: InfoHawk, Inc. | Contact: Rob Leathern, Cofounder and CEO | Date: 30 September 2026

1. Summary

We strongly support Ofcom's proposals. They are the most thoroughly considered set of proposals on fraudulent advertising we have seen. They layer measures carefully, involve third parties, and address newer risks such as advertisement generation tools that most regimes have not yet tackled. We particularly welcome the ad library measure (FAU E1 / FAS E1), including the requirement to cover all ads that UK users can encounter.

About us and declaration of interest. InfoHawk, Inc. finds and takes down harmful content and accounts, including fraudulent advertising, on behalf of businesses and consumers. Its co-founder and CEO, Rob Leathern, also sits on the board of CollectiveMetrics, a US-based non-profit working to measure the prevalence of safety and trust problems online. InfoHawk has a commercial interest in enhanced access to ad libraries, which we declare here. We believe the case set out below stands on its merits and would apply equally to any qualified organisation.

This response focuses on one point we think will decide how well ad libraries work in practice: access at scale for qualified organisations. Paragraph 5.95 of the consultation refers to an "enhanced-access API" as something providers may consider. We ask Ofcom to go further:

1. **Recognise the need.** Ofcom should acknowledge that high-volume access for qualified organisations is a distinct and necessary category of access.
2. **Require a process.** Providers should be expected to have a process for granting that access, based on objective criteria they set and publish.
3. **Don't over-specify.** Ofcom should avoid prescribing the criteria in detail, as the set of entities who are in a position to offer these kinds of services will evolve. Section 4 offers some suggestions providers could draw on.
4. **Link public restrictions to enhanced access.** Providers that tighten their limits (for example, anti-scraping) on the public library should also provide an enhanced route. Otherwise tighter public limits will reduce detection rather than protect it.

Sections 5 to 7 cover landing pages and cloaking, dates and reliability, reporting and feedback, the likely privacy objection, and measurement.

2. Who uses ad libraries

In our experience, ad library users fall into three broad groups:

- **Members of the public.** Someone sees an ad, suspects a scam, and wants to report it or warn others. They need a simple, reliable interface and a way to report.
- **Investigators without heavy technical resources.** Journalists and civil society researchers search for suspicious patterns, for example ads directing users to off-platform messaging apps, and bring them to public attention. They mostly use the web interface and generally do not need an API.

- **Organisations working at scale.** Some organisations, such as CollectiveMetrics, measure the prevalence of safety and trust problems; others, such as InfoHawk, find and take down bad content or accounts on behalf of businesses and consumers. Cybersecurity firms, for example, may search daily for ads or websites mentioning a well-known crypto wallet provider, including misspellings of its name, to find credential-phishing sites. Because ads are one of the attack vectors, these users need to pull large volumes of ads systematically on behalf of their clients.

The third group is currently the most poorly served and we number ourselves among this group. We are often blocked at every turn:

- Some ad libraries do not include all ads shown in a country, or do not cover the country at all. For example, a library may include ads shown in the UK but not those shown in other major markets.
- Some ad libraries omit click-through URLs altogether, or only include the URLs for a subset of the ads, or they may not provide enough information to simulate what a valid click looks like.
- Most limit API and scraping volumes so aggressively that systematic analysis is impractical.

The value this group creates goes beyond the ad platform itself. We recently found an ad using celebrity deepfakes to sell a “miracle cure” for Alzheimer’s. The same actors, who appeared to be operating from outside the UK, were also running the scheme through other ad networks and SMS. They relied on around 35 separate service providers, including domain registrars, hosting companies and video hosts. Organisations with scaled access can identify such networks and push them off not just ad platforms but the wider infrastructure that enables the scam.

3. Why enhanced access is needed (with limits)

The public library cannot serve high-volume users. Platforms limit scraping for good reasons (see below), so volume limits on the public library will stay tight. Without a separate route, the organisations best placed to find fraud at scale cannot do so. One platform has given InfoHawk access to its public ad library from specific IP addresses at 100 queries per second. This has let us extract ads at a useful scale, and shows the arrangement is not technically difficult. Similar requests to other platforms have not yet resulted in comparable access.

Unrestricted access would help bad actors. Fraudsters copy what they see working. Tracking ads over time shows which campaigns are succeeding: an ad that keeps reappearing is almost certainly performing. Advertising intelligence tools have used exactly this technique for some time now. Commercial “ad spy” tools already sell this kind of insight, and bad actors are among their customers. High-volume access therefore should not be open to everyone.

Access should not be limited by organisation type. Restricting access solely to academics or non-profits is not the answer either. Some such organisations lack the resources to use scaled access well, and the status could also be easy to fake: in some jurisdictions a non-profit can be registered very easily, and as an example we have seen bad actors register as a non-profit to obtain advertising grants intended for charities. Commercial organisations working on behalf of brands, such as brand protection and cybersecurity firms, could be among the most effective users of enhanced ad library access.

The balance depends on having both routes. Platforms will reasonably argue that exposing more data at volume helps bad actors. The answer is not to lock the library down; it is to pair tighter public limits with enhanced access for qualified organisations. If a casual user cannot retrieve even a handful of ads, and a security firm protecting brands has no route to retrieve thousands, the measure will fail.

It is a check on platforms meeting their obligations. Today it is near impossible to holistically understand if companies are compliant with even a fraction of their public policies and requirements, but the limited data we have inclines us to believe they are not doing so today. This kind of enhanced access is a key part of helping objective third parties have the ability to act in the public interest and help understand if these platforms are meeting their commitments and enforcing their policies.

4. How enhanced access could work

We suggest Ofcom require providers to have a process for enhanced access, but leave the design to them, consistent with market-based approaches being discussed in other jurisdictions. The suggestions below are illustrative, not a proposed specification.

4.1 Possible access arrangements

Providers could use any combination of:

- higher API rate limits for approved organisations;
- permitting approved organisations to exceed anti-scraping thresholds on public endpoints, for example from registered IP ranges;
- include more complete content and additional metadata for ads in a nonpublic feed (including for example information about the advertisers themselves that would allow the detection of coordinated networks of actors);
- potentially do not remove previously-taken down materials for those entities with ‘enhanced access’, since this is an important signal for finding repeat offenders and patterns;
- bulk export or change feeds; and
- impression- or revenue-weighted random samples of ads, up to an agreed limit. *These are essential for credible prevalence measurement (see Section 7).*

4.2 Suggested criteria providers could consider

Providers should set objective criteria and publish them. Examples they might draw on include:

- whether the organisation’s purpose is to detect, report, measure or disrupt fraudulent or harmful advertising, e.g. brand protection, security research or prevalence measurement (for example, use cases such as ad creative planning or generation should be disallowed);
- whether the organisation can demonstrate the value it will provide, such as a track record of reporting or published research;
- verification of the organisation’s identity, rather than reliance on its declared status;
- whether the organisation resells ad creative or campaign intelligence to advertisers. Brand protection users generally care who is running an ad, not what the ad shows, so a business model built on selling creative is a useful warning sign; and
- agreement to reasonable terms of use, with access that can be suspended if misused.

Ofcom's recommended trusted flaggers could reasonably be treated as eligible by default. Some procedural basics would also help: publishing the criteria, deciding applications in a reasonable time, and explaining refusals. That would make the criteria objective in practice and not just on paper.

4.3 What stays public

The public library should remain genuinely usable. Ofcom's reasons for making it available without registration are sound, and the first two user groups in Section 2 depend on it. Enhanced access should add to the public library, not replace it: providers should not be able to point to an enhanced route as a reason to make the public library unusable.

5. Other points on the ad library measure

5.1 Landing pages and cloaking

"Cloaking" is a major and growing problem: bad actors send platform reviewers and tools to one destination and real users to another. One tactic we see is fraudsters checking how the landing page was reached. If the visit did not come through the platform's click-tracking or redirect link, they hide the true destination. An ad library that simply shows the declared destination URL can therefore show a harmless page while users are sent somewhere else.

Some platforms handle this well: at least one major library replicates the platform's click-tracking link wrapper, apparently without inflating advertisers' click counts, so the actual destination can be observed. Other ad networks render URLs in ways that are very hard to reproduce, and we would expect a library built by such a network to show only the declared destination. We suggest Ofcom specify that, in the ad library, links to the landing page should replicate as closely as possible the technical path a real user's click follows. As more libraries become available, we expect fraudsters to try harder to game this.

5.2 Clarity about dates

In existing libraries it is often unclear whether a date means an ad started running, was running, or actually received impressions on that day. This confuses casual users and sometimes professionals too. We suggest the measure require clear labelling of when an ad started and stopped running and when it was served.

5.3 Payer information

We note that Ofcom's list omits Article 39's "who paid" field. In practice this field is of limited value today: it usually shows an entity name that tells users nothing. We would not prioritise copying it in its current form.

5.4 Reliability

We support reliability as a principle. Access problems come from two sources. Anti-scraping controls are legitimate. But libraries also often simply fail or slow down at busy times. Detailed standards are hard to set in advance. We suggest Ofcom build them up over time from real performance data, and consider periodic independent testing of availability and performance, at least for the largest services. This should cover both the public library and enhanced-access or trusted flagger routes.

6. Reporting, feedback and the privacy objection

6.1 Reporting from the library

We welcome the ability to report ads directly from the library. Providers should not be required to give outcome feedback to anonymous members of the public, since that could be weaponised by bad actors testing what gets removed. But providers should be able to systematically show that library reports are genuinely considered, not quietly discarded.

6.2 Feedback for trusted flaggers and other scaled reporters

For trusted flaggers and other organisations with access to the enhanced API and/or reporting at scale, two things matter:

- **A persistent tracking reference for each report.** Without one, an organisation that has submitted 50 reports cannot tell which have been reviewed.
- **Feedback on whether action was taken, and what action, within a reasonable time.** This lets reporters improve their own detection systems.

We understand Ofcom's concern about the burden of per-report feedback, and why the illegal harms approach was adapted here. However, reporting and review increasingly happen machine-to-machine. As they do, structured, automated feedback becomes far less costly, and we encourage Ofcom to revisit this over time. Whether removal reasons appear in the public library matters less than whether this feedback reaches trusted flaggers.

6.3 The likely privacy objection

Providers are likely to raise privacy concerns about publishing advertiser identities and previous account names. Advertising is public communication, and these concerns can be managed through clear disclosure at the point an ad is placed. For example, a sole trader should be told plainly that their name will appear publicly with the ad. There are already many mechanisms to manage company-level confidentiality in various jurisdictions globally that could be repurposed to counter any objections here.

There are edge cases where people may not realise they have become advertisers, such as individuals paying to promote a personal marketplace listing. One platform addressed this with prominent disclosures when the promotion was bought. Ofcom could consider making such disclosure an explicit expectation.

7. Measurement and industry collaboration

Although not asked directly in this section, we offer some observations relevant to performance targets (FAU C4) and to the value of enhanced access:

- **Prevalence needs weighted samples.** Credible prevalence figures require random samples weighted by impressions or revenue.
- **Time and reach must be measured together.** How long a fraudulent ad was live matters, but only alongside how many impressions it received. Ad delivery algorithms reward high engagement, so the most clickbait-style scam ads can gather very large reach very quickly.
- **Some headline metrics can mislead.** Proactive detection rates and user report rates are easy to present favourably. They should be heavily discounted unless paired with impression-weighted data.

- **Accuracy should be measured both ways.** Precision and recall, covering both false positives and false negatives, would be valuable for Ofcom, even if not meaningful to the public.
- **Revenue from removed ads could be reported.** We would welcome reporting of revenue earned from ads later removed, broken down by policy. Even an ad removed after a day can generate substantial revenue.
- **Fraudsters time their ads around review.** We have seen fraudsters infer that a platform's review takes about a day, run an ad for around 21 hours, remove it themselves, and relaunch it the next day. Takedown-time metrics alone will not capture this, so several elements such as how long an ad was up, when it was scheduled to run vs when it did run etc are all potentially useful signals.

Industry collaboration. We encourage Ofcom to make clear that providers can and should share information on emerging bad-actor trends with each other. Much of the most useful learning in this field comes from practitioners talking to one another.

8. Suggested drafting change

We suggest adding a provision along the following lines to FAU E1 / FAS E1. It deliberately leaves the criteria and access arrangements to providers:

> E1.x The provider should have and apply a process for granting enhanced access to the ad library to persons whose access is likely to assist in detecting, reporting, measuring or disrupting fraudulent advertising, including access at volumes materially greater than those available to the general public. The provider should set objective criteria for such access, publish a summary of them, and review them periodically. Where the provider restricts automated or high-volume access to the ad library, it should ensure that persons meeting its criteria have an effective means of obtaining such access.

We would also welcome further detail in the supporting text at paragraph 5.95, including illustrative examples of criteria and access arrangements such as those in Section 4.

9. Answers to the consultation questions (Volume 4, Section 5)

Do you agree with our proposal? Yes. We strongly support it, with the addition of an enhanced-access expectation as set out above.

Do you agree with the principles, functionalities and information categories? Broadly yes, particularly the inclusion of all ads, search functions, the API, advertiser information and landing pages. We suggest: requiring landing-page links that replicate a real user's click; clearer date labelling; building reliability standards over time with possible independent testing; and giving feedback to trusted flaggers rather than the public. We would not prioritise the Article 39 payer field in its current form.

Evidence on building or adjusting an ad library. We do not operate an ad library, so cannot give build costs. From the user side, our experience is set out in Sections 2 and 3: libraries that omit some countries or some ads; missing or partial click-through URLs, or too little information to simulate a real click; and rate limits that make systematic analysis impractical. Where one platform has granted IP-based access at 100 queries per second, large-scale extraction became practical, which suggests the technical cost of enhanced access is modest.

InfoHawk would welcome the opportunity to discuss these points further with Ofcom.