

Emerge Protocol Ltd

Company No. 16142850 | London, United Kingdom
www.emergedata.ai

Competition and Markets Authority — Digital Markets Unit
2 March 2026

Consultation Response: Data Portability Conduct Requirement

Google's General Search Services

1. About Emerge Protocol

Emerge Protocol Ltd (trading as Emerge; “Emerge”) is a UK-incorporated data personalisation company that operates a consent-driven data sharing platform. We enable consumers to connect their personal data from Google, Meta, Amazon, and other platforms through a secure Data Wallet, which is then shared with authorised partners via APIs and MCP integration. We are an active user of Google's Data Portability API (DPAPI) and operate a production export pipeline for three scopes: *chrome.history*, *myactivity.search*, and *myactivity.youtube*. We do not propose to claim anonymity for our views.

2. Executive Summary

Emerge strongly supports the proposed Data Portability CR and agrees with its stated aim. The CR will provide essential legal certainty. However, the proposed CR does not go far enough in one critical respect: **it does not address the absence of meaningful service level requirements around data delivery speed**. Without this, the most transformative use cases will remain commercially unviable.

Our submission makes three arguments:

1. **The modern consumer internet is frictionless.** 53% of mobile users abandon any experience taking longer than 3 seconds. Conversion drops 12% per additional second of delay. No consumer founder will build on infrastructure where data takes minutes or hours to arrive, and no user will tolerate it.
2. **Google's API is unreliable, not just slow.** Our production logs show that export duration does not correlate with data volume. A smaller export window can take longer than a larger one for the same user. This unpredictability makes the API unusable for any experience that requires a consistent user promise.
3. **The cost of improvement is minimal.** UK banks deliver Open Banking API responses in 325 milliseconds on average, processing 2 billion calls per month. Google already accesses this same user data in real time billions of times daily to serve its own products. The marginal cost of serving it to authorised third parties at comparable speed is trivial.

3. Only Frictionless Connections Will Get Adoption

3.1 Consumer expectations in 2026

Today's digital experiences are **entirely frictionless**. Instagram loads a personalised feed in under a second. Amazon completes a purchase instantly. Banking apps show balances in milliseconds. ChatGPT begins streaming responses immediately. This is the baseline every new consumer application must meet.

The data is unambiguous:

- **53% of mobile users abandon any experience taking longer than 3 seconds** (Google).
- **Conversion rates drop 12% for every additional second of delay** (industry benchmarks).
- **A 2-second delay results in abandonment rates of up to 87%** (UX research).
- **25% of users abandon a mobile app after a single use** if first value is not immediate (Statista/AppsFlyer).
- **Google's own research:** a 0.1-second improvement in mobile load time boosts retail conversion by 8.4%.

The DPAPI currently delivers data in **approximately 1 minute to over 100 minutes**. This is not a marginal miss—it is a fundamental incompatibility with how modern consumer products work. Even in the best cases we observed, delivery takes significantly longer than any consumer application can absorb. In typical cases, users are waiting minutes or more for data that every viable consumer experience requires in seconds. No product category we engaged with could build around this constraint.

3.2 Use cases that depend on speed

We have engaged with 16 companies across 7 distinct use case categories (see Appendix A). These span AI companions, dating, rewards, travel personalisation, e-commerce intent, edtech, and fintech. Collectively, these companies represent between 1.5 and 2 million addressable UK users—impact that could be unlocked within the first six months of a functioning API, based on our direct conversations with their founders and CPOs. Six of the seven categories require near-real-time data delivery and are blocked or severely degraded by the current API performance. A summary:

- **AI companions and personal AI (3 companies):** Require real-time. Consumers win: AI that knows you from day one without weeks of interrogation. Instead of answering dozens of onboarding questions or starting from a blank slate, users get an AI assistant that immediately understands their interests, routines, and preferences—saving time and delivering genuinely useful recommendations from the very first interaction.
- **Dating and matchmaking (2 companies):** Require real-time. Consumers win: matching on real behavioural signals rather than self-reported profiles. Users are paired based on what they actually do—shared interests evident from browsing and search behaviour—rather than aspirational self-descriptions, leading to higher-quality matches and reduced time wasted on incompatible connections.
- **Rewards and data monetisation (4 companies):** Require fast delivery. Consumers win: fair compensation for data they already generate. Users currently create enormous value for platforms through their browsing and purchase behaviour but receive nothing in return. Data portability enables consumers to share this data on their own terms and receive tangible rewards—cashback, discounts, or direct payments—creating an entirely new value exchange.

- **Travel personalisation (2 companies):** Require real-time for in-session use. One company confirmed >20–30 seconds blocks their most valuable use cases. Consumers win: instant personalisation that eliminates the overwhelming catalogue experience. Instead of scrolling through thousands of generic hotel and flight options, travellers see results tailored to their actual preferences, budgets, and travel patterns—reducing decision fatigue and helping them find better deals faster.
- **E-commerce intent (2 companies):** Require real-time. Consumers win: proactive, relevant offers at the moment of decision rather than intrusive retargeting after the fact. When a consumer is actively researching a purchase, services can surface genuinely useful comparisons, price alerts, or alternatives—delivering value precisely when the consumer wants it, not days later through disruptive advertising.
- **Fintech cold-start (2 companies):** Require real-time for onboarding. Consumers win: personalised financial tools from day one. New fintech apps can immediately understand a user’s financial interests and spending context, enabling tailored budgeting suggestions, relevant savings products, and appropriate investment options without requiring months of transaction history to build a profile.
- **EdTech (1 company):** Moderate urgency. Consumers win: recommendations based on actual learning behaviour rather than generic course catalogues. Learners receive content suggestions aligned with what they have genuinely explored and engaged with, helping them progress more efficiently and discover relevant courses they would not have found through traditional search.

The full table of use cases, addressable users, and consumer benefits is provided as **Appendix A** alongside this submission.

4. Evidence: The API is Unreliable, Not Just Slow

We ran a controlled benchmark across our production pipeline comparing three export windows: 3 months, 13 months, and all-time. The API offers the ability to select an export window, but our evidence shows this does not translate into predictable or proportional delivery times. The core finding is that the API’s behaviour is fundamentally inconsistent.

4.1 Export duration does not correlate with data volume

Below are illustrative examples from our production logs. We have observed this inconsistency across many users—these are representative, not outliers:

User	Larger Window	Duration	Smaller Window	Duration
User A	13 months	1m 19s	3 months	4m 10s
User B	13 months	4m 11s	3 months	6m 43s

Table 1: Illustrative examples — larger export windows completing faster than smaller ones for the same user. Full Google Job IDs available to the CMA on request. This pattern is consistent across our user base.

The distribution of export times across all completed jobs confirms this:

Export Window	Range	Typical Size
3 months	1m – 6m 44s+	~2 MB
13 months	~1m – 16m 11s+	~9 MB
All time	1m – 102m 35s	~9 MB

Table 2: Export duration ranges overlap significantly across all three windows despite vastly different data volumes. The bottleneck is queuing and scheduling, not data volume.

This evidence demonstrates that the latency is an **engineering choice, not a technical constraint**. If export duration were driven by data volume, smaller windows would consistently be faster. They are not. The API's architecture appears to process jobs through a queue that introduces arbitrary delay regardless of request size, making even small, targeted requests completely unpredictable.

We anticipate that Google may argue that third parties seeking faster results should simply request a shorter export timeframe. Our evidence directly refutes this. As Table 1 demonstrates, a 3-month export can take **longer** than a 13-month export for the same user. Reducing the export window does not produce proportionally faster delivery because the bottleneck is not data retrieval—it is the queue itself. A developer cannot reliably promise a user “this will take 10 seconds” regardless of which timeframe they select. The unpredictability is the problem, and it persists across all window sizes. No combination of parameters available to third-party developers can overcome a queue-based architecture that introduces non-deterministic delay.

4.2 “All-time” exports are truncated

We also observed that “all-time” exports consistently returned only approximately 12 months of records, even where user history spans significantly longer. This is a data completeness concern relevant to Interpretative Note 3(a)(iii).

5. The Consent Flow Discourages Data Portability

The consultation notes drop-out rates of **over 70%** during Google's consent flow (paragraph 1.17). Our own internal UX research confirms this. We have conducted structured testing with over 20 participants, observing their experience of the consent flow in real time. In every session, we see two consistent reactions:

- **A “wow” moment** when users see the richness of data available to them and understand the value of sharing it. This is precisely the reaction that data portability should produce. It is the moment a user realises they have agency over their own digital footprint. We actively design our onboarding to amplify this moment, because an empowered user who understands what they are sharing—and why—is the best possible foundation for a healthy data ecosystem. This is what good consent looks like.
- **Significant concern and hesitation** at the Google consent screen that displays an orange exclamation mark warning triangle. This visual element is critically important to understand in context. Google does not deploy this iconography casually. Across the entire Google product ecosystem—Gmail, Chrome, Android, Google Drive, Google Pay—this orange warning triangle is reserved exclusively for genuinely dangerous situations: phishing attempts, malware downloads, compromised accounts, and unsafe websites. It is the visual language Google uses to tell users “stop, something here could harm you.” The fact that Google deploys this same alarm signal when a user exercises their legal right to port their own data is extraordinary. It treats a legitimate, CMA-supported action as though it were a security threat on par with a phishing attack.

The warning iconography must be removed. Presenting the same visual signal when a user exercises their right to port their own data is not an “appropriate balance” of ease and security (Interpretative Note 3(g))—it is a design choice that actively undermines the purpose of the CR. The distinction between informed consent and manufactured fear is important.

Informed consent means giving users clear, neutral information about what data will be shared and with whom. Manufactured fear means deploying visual threat cues that exploit years of security conditioning to discourage a perfectly safe action. Our testing consistently shows that the current flow does the latter. Users who moments earlier expressed genuine excitement about their data’s potential are visibly deterred by a warning pattern they have been trained to associate with danger.

We strongly urge the CMA to require the removal of warning iconography from the consent flow and to benchmark Google’s consent UX against its own patterns for equivalent user actions. When a user grants a third-party app access to their Google account via OAuth—a flow Google facilitates billions of times annually—no such warning appears. The warning is reserved specifically for data portability. This is not a neutral design choice: it is the deployment of Google’s most alarming visual language against the precise user action the CMA seeks to encourage.

6. Response to Google’s Submissions

Google submitted (paragraph 1.20) that it has “neither the incentive, nor the ability, to degrade the performance of the API” and that doing so “would damage Google’s reputation.” We respectfully submit this argument deserves limited weight. Google’s reputation is not built on the DPAPI—it is a regulatory compliance tool with a tiny user base. More importantly, the question is not whether Google would *overtly* degrade the API, but whether it has any incentive to make it genuinely fast enough to enable third-party innovation that competes with its own services. The current performance—slow, unpredictable, and inconsistent—is entirely consistent with **an absence of incentive to improve**.

7. The Cost of Improvement is Minimal

Google routinely serves billions of personalised search results per day in milliseconds, drawing on the same user data the DPAPI exports. The question is not whether Google can deliver this data quickly—it demonstrably can—but whether it chooses to.

When the CMA mandated Open Banking, UK banks built real-time APIs *from scratch on legacy infrastructure*. The ecosystem now processes over 2 billion calls per month at an average response time of **325 milliseconds** (January 2026). Google already has the API in place, the data indexed, and volumes are modest (single-digit MBs per user). Our evidence (Section 4) shows the bottleneck is in queuing, not processing—a routine engineering optimisation. The CMA’s own cost estimate of up to £1 million per year (paragraph 5.13) easily covers this.

8. Recommendations

4. **Introduce a baseline performance requirement:** For exports up to 6 months, data should be delivered within 10 seconds. This is within Google’s capabilities, consistent with Open Banking, and would unlock the next generation of consumer applications.
5. **Require transparency on response times:** Expand compliance metrics (paragraph 4.22) to include average and 95th percentile delivery times by export window, published in the non-confidential compliance report. Transparency drove continuous improvement in Open Banking.
6. **Address the consent flow:** Require Google to benchmark its consent flow against its own UX standards and remove design elements inconsistent with how it presents routine actions elsewhere.

7. **Signal direction:** Even if a binding SLR is premature, state clearly that near-real-time delivery is within expectations for “effective” portability and will be formally required if not achieved. This gives founders and investors the confidence they need.

9. Responses to Consultation Questions

(a) Do you agree with the aim?

Yes. The CR will unlock innovation that cannot be predicted in advance—as the App Store and Open Banking both demonstrated. However, effectiveness is contingent on the API performing at a level compatible with modern consumer applications.

(b) Would it result in the benefits identified?

Yes, but **only if operational standards support adoption**. Without near-real-time delivery, the CR risks becoming a paper right. With it, benefits will substantially exceed the CMA’s conservative estimates. The £4 billion Open Banking comparison is apt—but only if equivalent performance standards apply.

(c)–(d) Do you agree with the Interpretative Notes?

Broadly yes, but the CMA should go further. **Google’s current implementation should not be the baseline**. Terms like “reasonable business needs” need specificity to prevent Google interpreting them to justify preserving the status quo.

(e) Compliance and monitoring?

Sensible. Add API response time metrics. Signal willingness to escalate as usage grows. Keep the consent flow under active review.

10. Conclusion

The proposed CR is a significant step. But legal certainty alone is not sufficient. Millions of UK consumers are currently unable to benefit from personalised services because the infrastructure that would enable those services is too slow to support them. This is not foregone innovation—it is existing consumer harm, happening now. The transformative potential of data portability requires that **data moves at the speed of the modern internet**. Our production evidence shows the API is unreliable, not just slow. The Open Banking precedent shows regulating for performance works. And 16 companies across 7 use case categories—representing 1.5 to 2 million UK users—are ready to build within months of a functioning API, if the infrastructure performs.

Emerge is willing to provide further evidence, including full production benchmark logs, and to make our founder network available for further discussions with the CMA team.

Submitted by:



Chief Executive Officer & Founder, Emerge Protocol Ltd

2 March 2026

Appendix A: Use Case Market Need

The detailed use case table documenting 16 companies across 7 categories, their addressable user bases, real-time requirements, and specific consumer benefits is provided as a **separate Excel file** alongside this submission. The table demonstrates that 6 of 7 use case categories—representing the vast majority of addressable users—require near-real-time data delivery to support viable consumer products.

Note: Company names have been withheld to protect commercial confidentiality. EmERGE can facilitate introductions to any of these founders upon the CMA's request.