Response to the public consultation on Guidelines 02/2026 on Anonymisation

Guidelines	Guidelines 02/2026 on Anonymisation, version 1.0, adopted 7 July 2026
Respondent	Hugh Ndhlovu, trading as NooSight
Type of respondent	Small developer of privacy-focused consumer software, established in the United Kingdom, which may in future offer services to users in the European Economic Area
Date	15 September 2026

1. Summary

We welcome the draft Guidelines, and in particular the assessment of anonymity from each relevant entity's perspective under the contextual approach, and the three criteria of No Record Isolation, No Linkage and No Inference (section 3.4). We design consumer software in which personal data stays on the user's device. We are designing an opt-in way for users to contribute reduced reports for publication as aggregate statistics; it is not built, and no contribution has been sent.

For designs of that kind the Guidelines raise three practical questions that we believe many developers face, and on which a worked example in the final Guidelines would help. We set out a hypothetical design below, our own view on each question with our reasons, and the question on which we would value the Board's view.

2. The hypothetical design

A consumer app runs on the user's phone and helps people keep a record of their dealings with businesses. Personal data stays on the phone, except for the reduced contributions described here. Users may opt in to contribute reports about how businesses handle a routine consumer request, so that a central service can publish statistics about each business.

A contribution contains only:

1. an identifier for the business (a company, not a person);
2. a broad category of service, from a short fixed list;
3. an outcome, from a short fixed list (for example, "request acknowledged" or "no response");
4. the month and year.

It contains no free text, no account or device identifier, no location, and no exact time. The reduction happens on the device before anything is sent.

3. The controls under consideration

1. **Abuse prevention without a recognisable identifier.** To stop one person flooding the statistics, each phone may make a limited number of contributions in a period. Rather than a

device key the central service would recognise, the design uses one-time allowance tokens of the kind described in the IETF Privacy Pass documents: a separate issuing function grants tokens, and the token is built so that the service which receives a contribution cannot link it to the issuance, or to any other contribution. The issuing function keeps a short-lived record to enforce the limit. That record is treated as personal data, kept apart from the contributions, and deleted on a short fixed schedule.

2. **A short record-level window.** The central service holds individual contributions for a short fixed period to check for fraud (for example, bursts of identical reports), then adds them to aggregate counts and deletes the individual records.
3. **Publication thresholds.** A figure is published only when it rests on a minimum number of contributions; combinations of fields that would fall below the minimum are suppressed, and users of the published statistics cannot filter across fields to recover small cells.
4. **Sensitive categories and businesses excluded at source.** Where the category of service, or the business itself, could reveal special category data (for example, a health clinic or a religious organisation), no contribution is made by default. The check runs on the device against a fixed list, and looks at the business as well as its category, because a business's identity can reveal what the category hides.
5. **Network details not kept.** The device's IP address and other connection details reach the receiving service in transit, as with any request. They are not written to logs or stored with a contribution, and the issuing and receiving functions do not share network logs. Contributions may be delayed on the device by a random interval, so that the time a token is issued cannot be matched to the time a contribution arrives.
6. **Reassessment.** Any change to the fields, the thresholds or the token design triggers a fresh assessment against the three criteria before it takes effect, and a security incident affecting the issuing function is handled as a personal data breach.

4. Question 1: linkage within one organisation

OUR VIEW. Assessed from the perspective of the receiving service, the aggregate dataset it keeps after the window described in section 5 meets the No Linkage criterion, even though the issuing function holds personal data for a short time. The design aims to remove the means of linkage, not only to forbid it by policy: the token carries nothing that joins a contribution to its issuance, network details are not kept, a random delay breaks matching by time, and the issuing record is deleted quickly. We treat the issuing record as personal data throughout.

OUR REASONS. The contextual approach asks whether the entity holding the data has means reasonably likely to be used to identify a person (sections 2.4.2 and 3.5.2). Where the link is removed by the design of the token, together with the handling of network details and timing, rather than held behind an internal access rule, we consider that the receiving service has no means reasonably likely to be used. We accept that this depends on those measures working as designed, which is why control 6 requires reassessment on any change.

> **OUR QUESTION.** Where both functions are operated by the same controller, does the Board consider that the contribution dataset can be anonymous for that controller, provided the unlinkability rests on the design of the token rather than on organisational separation? If

> organisational separation is also expected, what degree of separation would the Board regard as sufficient?

5. Question 2: the short record-level window

OUR VIEW. During the fraud-check window, individual contributions should be treated as personal data, and the processing given a lawful basis and security to match, even though each contains only four coarse fields. They become anonymous for the service once aggregated and deleted at record level.

OUR REASONS. This is the cautious reading. A single record, taken with outside knowledge (a user who tells others they complained to a business in a given month), could in some cases be isolated. We would rather carry the obligations for a short window than rely on an argument that four coarse fields can never isolate a record.

> **OUR QUESTION.** Does the Board agree that this cautious reading is appropriate? Or would the Board regard record-level contributions of this coarseness, holding no identifier and received through unlinkable tokens, as capable of being anonymous from the moment of receipt?

6. Question 3: inference from the category of service

OUR VIEW. Excluding special category services at source is necessary but not always enough. Some categories are not special category data but still reveal sensitive facts about a person, such as debt advice or gambling services. For those, a higher publication threshold, or not publishing below a wider grouping, is a proportionate step.

> **OUR QUESTION.** Would the Board expect the No Inference criterion to be met by a fixed exclusion list for special category services plus uniform thresholds? Or does it expect thresholds that rise with how sensitive a category is, and if so, would the final Guidelines give an example?

7. Closing

We would welcome a worked example in the final Guidelines covering opt-in contributions from consumer devices to a published statistical service, because this pattern lets people share useful information about how businesses treat them without giving up their privacy. We are happy for this response to be published.