BeeWii

DOCUMENT BEEWII

Comments on Guidelines 02/2026 on Anonymisation

Référence	BW-EDPB-ANON-2026-COM-01
Version	1.0
Date	Oct 12, 2026
Statut	Validé
Classification	Public

Table of Contents

About BeeWii....2

1. Free text needs its own guidance....3

2. Automated detection always leaves something behind....4

3. Prompts pseudonymised before they reach an AI provider....5

4. Where the anonymising step runs matters....6

Summary of requests....7

Sources....8

About BeeWii

BeeWii is a small French software company. We build open-source tools that let organisations use AI while keeping control of their data. We welcome these guidelines. The three criteria and the two approaches give practitioners a structure they did not have since Opinion 05/2014.

Our comments come from tools we build and test, not from legal practice:

- go-anon, a library that pseudonymises or redacts Word, OpenDocument, PDF and CSV files, then runs a second detection pass on its own output to check for leaks.
- indecis-ner, a 7.7-million-parameter model that finds direct identifiers, quasi-identifiers and Article 9 data in French prose.
- indecis-caviardage, a demonstrator that combines the two. It runs on the user's machine or entirely in the browser, and the document is never sent anywhere.
- Xolo, an AI gateway that sits between users and LLM providers. One of its plugins pseudonymises prompts and attachments before they leave the organisation, and restores the original values in the response.

All four handle free text: documents, emails, prompts. Our comments focus on that case, because we think the guidelines say little about it.

1. Free text needs its own guidance

The technical framework in section 3 is written for tables. Every example is a table or an aggregate. Paragraph 85 defines dimensionality as "the number of columns in a table". Yet much of the data organisations want to share or reuse is prose. Court decisions, case files, HR reports, medical letters, emails and, more and more, prompts sent to AI systems.

Applying the three criteria to prose raises questions the draft does not answer.

- **What is a record?** Paragraph 57 says a record may span "multiple lines in a table". In a document, what relates to one person is spread over sentences, pronouns and indirect references ("her brother", "the new head of the Dijon branch"). Nothing marks where one person's record ends.
- **Which attributes count?** A table has a fixed set of columns. Prose does not. Any sentence can add an attribute, and the author decides which ones, not a schema. Dimensionality and resolution (paragraph 85) are then hard to measure.
- **Narrative quasi-identifiers.** Once names are gone, a story can still single someone out. "The only nurse in the department who came back from sick leave in March" isolates a record with no identifier in it. This is why our model also tags occupations, relatives and life events, not only names, dates and places.

We would ask the EDPB to add at least one worked example on a free-text document. It would help to say how to delimit a record in prose, and whether removing direct identifiers can ever meet the simplified approach for documents. Our experience suggests it rarely can, and that the contextual approach is the realistic route for most document publication.

2. Automated detection always leaves something behind

Tools that find and mask personal data in text make mistakes, and the rate can be measured. In our tests on French prose, about 2% of personal passages stayed visible at the default setting, and about 1% one step more cautious. Each step up also masks more words that needed no masking. Measured on other text, the numbers would differ, but they will not reach zero.

Paragraph 87 says a technique succeeds "even if it only succeeds against a single individual". Read with a 1% miss rate, any corpus of more than a few dozen documents will almost certainly contain a visible quasi-identifier. We agree with that reading. Market practice does not follow it. Many products sold as "anonymisation" are detection-and-masking tools, and their output is pseudonymised or redacted data.

We would ask the EDPB to:

- state plainly that masking detected entities in text is not, by itself, anonymisation, and that paragraph 40 applies to the vendors of these tools as well as to controllers;
- say what a controller should document when an automated tool is part of the process (paragraph 41). In our view that means the measured miss rate on comparable text, the test set used, the setting chosen, and whether a person reviewed the output;
- say whether a second, independent detection pass on the output and a human review count as a "safety margin" in the sense of paragraph 88.

Our own tools write out the full list of what they replaced, with the type and context of each item. We publish the code and the evaluation method so that others can check the numbers.

3. Prompts pseudonymised before they reach an AI provider

This case is now common and the guidelines do not cover it. An organisation routes its staff's requests to a hosted language model through a gateway. The gateway replaces names, addresses and other identifiers with placeholders such as `PERSON_1`, sends the prompt, then puts the original values back into the answer. The table that maps placeholders to values never leaves the organisation.

We read the guidelines as saying that these prompts remain personal data for the provider in most cases.

- When the provider acts as a processor, the controller's perspective applies to it (paragraph 15, Example 4). The controller holds the mapping, so the data is personal for both.
- When the provider acts on its own account, its own perspective applies, but the criteria are hard to meet. A prompt is often a single person's situation told in full, so No Record Isolation fails. The provider also knows which customer account, and often which user, sent the request, which opens linkage.
- The recipient is itself a language model. Paragraphs 82 and 92 note that AI lowers the cost of inference and linkage. Here the data is handed directly to the kind of system those paragraphs describe.

We built such a gateway and we still think it is worth having. It reduces what the provider receives and it limits the damage of a leak on the provider's side. But it is a pseudonymisation measure under Article 4(5) and Article 32, not anonymisation.

We would ask the EDPB to add an example on this case, and to confirm that offers presenting AI use as "anonymous" because prompts are masked fall under paragraph 40.

4. Where the anonymising step runs matters

Paragraph 38 reminds controllers that anonymisation is itself processing. In practice, many text anonymisation services run in the cloud. To anonymise a document, the controller first sends the original, with every identifier in it, to a third party. That adds a processor, sometimes a transfer outside the EU, and one more place where the original and the placeholder mapping can leak.

This need not be the case. Our demonstrator processes a ten-page document in 1.3 seconds on one laptop processor core, and the browser version runs without any server. The original and the mapping stay on the user's machine and are gone when the page closes.

We would suggest that section 2.5.2 mention where the anonymisation runs and who can access the originals and the mapping during the process. Local processing is one way to reduce that exposure. It fits with paragraph 39, which treats immediate anonymisation and deletion of the original as a positive factor.

Summary of requests

1. Add a worked example on a free-text document, and explain how to delimit a record and count attributes in prose.
2. State that masking detected entities in text is not anonymisation by itself, and apply paragraph 40 to the tools sold for it.
3. Say what to document when an automated detection tool is used, including its measured miss rate, and whether a second detection pass and human review count as safety margins.
4. Add an example on prompts pseudonymised by a gateway before they reach an AI provider.
5. Mention, in section 2.5.2, where the anonymising step runs and who can access the originals and the mapping.

We are happy to share our test sets and measurements with the EDPB.

Contact: BeeWii – bonjour@beewii.fr.

Sources

- EDPB Guidelines 02/2026 on Anonymisation, version 1.0 for public consultation, adopted 7 July 2026
- indecis-caviardage, README: measured miss rates and processing times
- indecis-ner, README: evaluation method and its limits