INSIGHT · GDPR & DPIA

Why most DPIAs fail before the first interview

By Shalini Jhunjhunwala · Aequita Advisory · 8 min read

A surprising number of DPIAs are dead before the first interview ends. Not because the regulation is unclear, and not because the assessor is inexperienced. The assessment starts before anyone has done the homework that makes it possible to do well.

I have run DPIAs in-house, on the consulting side, and as part of compliance programmes across regulated industries. The pattern repeats. A two-hour kickoff meeting that should have produced an initial risk picture instead produces a list of questions for a second meeting, which produces a list for a third. The DPIA gets done eventually, but the cost has tripled and the document everyone signs feels obligatory rather than illuminating.

Here is what actually goes wrong, and what to do about it.

The diagnosis: the assessment starts too early

GDPR Article 35 sets out when a DPIA is required and what it must cover. What it does not say, because it cannot, is when an organisation is *ready* for one. Most teams conflate the two. Once it is decided that a DPIA is needed, the assessment kicks off immediately, and the first meeting becomes the place where everyone discovers what they do not know.

The fundamental confusion is this: a DPIA is a **risk assessment**, not a **fact-finding mission**. The two work in opposite directions. A fact-finding mission collects information. A risk assessment evaluates information against a framework of harms and mitigations. Trying to do both in the same room with the same people at the same time produces an exhausting meeting and a thin output.

The fix is upstream. Before the first interview, the product team should have answered, in writing, about a dozen questions whose answers determine the shape of the assessment. When those answers exist, the kickoff becomes what it should be: a conversation about risk, not a conversation about facts.

The twelve questions that make the difference

Across the DPIAs I have been involved in, the same twelve questions account for most of the productive ground covered in early meetings. Sending them to the product team before the kickoff, and asking for written answers, typically cuts the first interview from two hours to forty-five minutes, and dramatically improves the quality of the eventual document.

The questions fall into four clusters:

- **What is being processed.** Plain-English description, categories of personal data, special-category data under Article 9, treatment of vulnerable groups.
- **Where the data goes.** Sub-processors and third parties, transfers outside the EU/EEA, retention and deletion mechanisms.
- **How decisions are made.** Automated decisions affecting users, machine-learning models involved, human review and override mechanisms.
- **What could go wrong.** Worst-case failure modes from the engineers' perspective, existing preventative and detective controls.

None of these questions are surprising. None of them require legal expertise to answer. All of them are routinely the things product and engineering teams cannot fluently answer in a meeting unless they have been asked beforehand.

What this changes about the assessment itself

When the twelve questions are answered in advance, three things shift.

The kickoff becomes a working session, not a discovery session. Time is spent stress-testing the team's understanding of their own system, identifying gaps, and probing assumptions, rather than soliciting basic facts.

The risk picture forms faster. With concrete facts in front of you, the shape of the risk becomes visible in the first hour. The patterns that matter, such as special-category data flowing to non-EU sub-processors, automated decisions without meaningful human review, or retention periods without deletion mechanisms, surface immediately.

The document at the end says something. A DPIA that emerges from a well-prepared kickoff reads like genuine analysis. A DPIA that emerges from a poorly prepared kickoff reads like a compliance artefact, with the box checked and no real insight. A supervisory authority can tell the difference.

One caveat about vague answers

A common reaction to the pre-interview questions is that some answers will be partial or uncertain. That is fine. In fact, a vague answer to a sharp question is itself useful diagnostic information.

If the team cannot say what sub-processors are touching the data, the DPIA needs to make that gap visible. If retention periods are not actually enforced anywhere, that is a finding. If automated decisions are made without a documented override mechanism, the mitigation section writes itself.

The point of the pre-interview is not to grade the product team. It is to see the shape of the work before anyone walks into the room.

Closing thought

Most DPIAs that feel performative, bureaucratic, expensive, or unsatisfying are not performative because of the regulation. They are performative because the assessment started before the assessment was possible. The fix is structural: do the upstream work, and the downstream work becomes useful again.

If you are running, scoping, or considering a DPIA, or if your team is the one being asked to produce one, the right question is not *have we started*. It is *have we done the homework*.

Aequita Advisory
Privacy & AI Compliance · EU practice
hello@aequita.eu
Shalini Jhunjhunwala, Founder