The DPIA guide for AI.

A Data Protection Impact Assessment is the paperwork that proves you thought before you processed. Under UK and EU GDPR it is legally required for most AI use involving personal data — and it is the first document a regulator, an insurer, or your biggest client will ask to see.

Written for the small firm putting corporate data — and especially *client* data — through AI tools. Includes the screening test, the AI risk catalogue, a fill-in template, and a worked example.

document	dpia-guide-for-ai
author	peter bassill
version	1.0 · 2026-07
locale	en_GB · UTC
licence	CC BY 4.0

§ 01 / WHAT AND WHEN

Thinking on paper, before the data moves

A DPIA is a structured written assessment of what a piece of processing does to people: what data, whose, why, what could go wrong, and what you did about it. Article 35 of both the UK GDPR and the EU GDPR makes one mandatory wherever processing is "*likely to result in a high risk to the rights and freedoms of natural persons*" — and the regulators on both sides have been explicit that AI usually clears that bar.

The reasoning is not anti-AI paranoia. AI tools concentrate several classic risk factors at once: they are novel technology, they often evaluate or profile people, they move data to third parties, they can produce decisions or content about individuals, and their failure modes — leakage, inaccuracy stated with confidence, unintended retention — are exactly the harms Article 35 was written for. The ICO's AI guidance treats a DPIA as the default expectation for AI involving personal data, not an edge case.

THE TWO SENTENCES THAT REFRAME THE CHORE

A DPIA is not compliance theatre — it is the AI risk assessment you should be doing anyway, written in the format that counts as a legal defence. Done honestly it takes an afternoon, kills bad ideas cheaply while they are still ideas, and answers in advance the exact questions a client's procurement team will send you.

When is one legally required?

Screen every new AI use against these triggers. One strong trigger, or two weaker ones together, means do the DPIA:

- **Innovative technology** — AI and machine learning are named examples in the ICO's high-risk indicators and the EDPB's criteria. Almost every tool in this kit's scope qualifies.
- **Evaluation, scoring or profiling** — AI that assesses people: CV screening, credit-worthiness, performance analysis, customer segmentation, sentiment analysis on client calls.
- **Automated decisions with real effects** — anything where the machine's output materially affects a person: pricing, shortlisting, service eligibility.
- **Special category or criminal-offence data** — health, biometrics, politics, religion, sexuality — at any meaningful scale. A transcription tool in a clinic hits this on day one.
- **Large-scale processing or systematic monitoring** — AI across a whole customer base, or watching behaviour over time.
- **Data about vulnerable people** — children, patients, employees (the ICO treats the employment power imbalance as a vulnerability factor).
- **Matching or combining datasets** — AI that joins sources which were collected separately, for different purposes.

Practical rule for a small firm: if an AI tool touches personal data and does more than reword text you could have retyped yourself, screen it. The screening record — even where the answer is "no DPIA needed" — is itself evidence of accountability, and takes ten minutes.

§ 02 / TWO REGIMES

UK GDPR and EU GDPR: same skeleton, diverging edges

If you only serve UK clients, the ICO's rules are your world. Sell into the EU, process EU residents' data, or serve clients who do, and both regimes apply at once. The good news: one well-built DPIA satisfies both, because Article 35 is near-identical. The differences sit at the edges.

	UK GDPR (ICO)	EU GDPR (EDPB / YOUR LEAD AUTHORITY)
When required	Art 35 + the ICO's published list of high-risk examples — which names innovative technology such as AI explicitly.	Art 35 + EDPB's nine criteria (WP248): two or more usually means a DPIA. Member-state authorities publish their own mandatory lists on top.
Automated decisions	Reshaped by the Data (Use and Access) Act 2025: solely-automated decisions with significant effects are more broadly permitted, provided safeguards — telling people, human review on request, contest rights. Stricter rules persist for special category data.	Art 22 remains: solely-automated decisions with legal or similarly significant effects are prohibited except narrow gateways (contract, consent, law), always with safeguards.
Residual high risk	Prior consultation with the ICO before processing starts (Art 36).	Prior consultation with the supervisory authority (Art 36).
AI-law overlay	No AI act; the ICO's AI guidance and the security duty in Art 32 do the work.	The EU AI Act adds a separate Fundamental Rights Impact Assessment (FRIA) for some high-risk deployers — distinct from, but overlapping with, the DPIA. High-risk obligations now bite from late 2027 after the 2026 postponements.
Transfers	UK IDTA / UK addendum to EU SCCs; UK–US Data Bridge for certified recipients.	EU SCCs; EU–US Data Privacy Framework for certified recipients.

WHAT EVERY DPIA MUST CONTAIN — BOTH REGIMES, ART 35(7)

- **A systematic description** of the processing and its purposes — for AI: the tool, the model, the data flows, who the provider is, where the data goes.
- **Necessity and proportionality** — why this processing, this much data, this tool; what less-intrusive alternative you considered.
- **The risks** to individuals' rights and freedoms — assessed for likelihood and severity, from their point of view, not yours.
- **The measures** addressing those risks — controls, safeguards, and the residual risk that remains after them.

The template in §05 follows exactly this structure, so completing it *is* compliance with the content requirement.

§ 03 / THE RISKS

The AI risk catalogue — steal these for section 4 of your DPIA

The hardest page of any DPIA is the risk table — a blank box invites panic or complacency. These are the risks that actually recur when small firms put personal data through AI. Copy what applies; strike what doesn't.

RISK	WHAT IT LOOKS LIKE IN PRACTICE	TYPICAL MITIGATIONS
Training-data absorption	Data pasted into a consumer-tier tool trains the provider's models; it cannot be recalled.	Business tiers; training off, evidenced; three-bin rule (doc 03).
Unintended disclosure	One user's data surfaces to another — shared workspaces, history features, provider incident.	Access controls, SSO, workspace config review, breach-notice terms.
Confident inaccuracy	The model states wrong facts about a person; the output enters a report or decision.	Human review before use (R.07); check against source records; correction route.
Unfair or biased outputs	Scoring or screening that treats groups differently — invisible in any single case.	AI advisory-only in decisions about people; sample-test outputs; named human decision-maker.
Automated-decision breach	The "human in the loop" is a rubber stamp — solely automated in law, without safeguards.	Real review with authority to disagree; tell people; contest routes (UK DUAA / EU Art 22).
Unlawful transfer	Data processed outside the UK/EU without a valid transfer mechanism — a common vendor default.	Check processing location; IDTA/SCCs/adequacy; regional hosting where offered.
Retention drift	Prompts and histories persist long after the source data should have been deleted.	Retention settings; scheduled deletion; AI tools inside deletion procedures.
Agent overreach	An agent acts on personal data beyond its task — reads the whole mailbox for one question.	Least privilege, scoped credentials, approval gates, logging (doc 04).
Sub-processor sprawl	Your vendor's vendors — model hosts, transcription, analytics — each a link you answer for.	Read the sub-processor list; DPA flow-down; four supplier questions (doc 01).
Rights-request failure	A subject access or deletion request arrives; nobody can find what the AI tools hold.	Map AI-held personal data before go-live; verify export/delete controls work.

Scoring: rate each applicable risk for **likelihood** (rare / possible / likely) and **severity** (minimal / significant / severe) *before* and *after* mitigations. Severity is judged from the individual's perspective — "embarrassing for us" is not the test; "harmful to them" is.

§ 04 / CLIENT DATA

When the data is your client's, the stakes double

Everything so far assumed the personal data was yours to decide about. For agencies, accountants, MSPs, recruiters, consultants — anyone processing *client* data — it usually isn't. Before any client data goes near an AI tool, answer one question precisely: **whose processing is this?**

The three positions, and what each demands

YOU ARE…	WHEN	WHAT THAT MEANS FOR AI USE
Processor	You handle their data on their instructions — payroll, support tickets, campaign lists.	Your contract (the DPA) governs. Adding an AI tool = adding a sub-processor: most DPAs require notice or consent *first*. The DPIA duty is primarily the client's — but they cannot do it without your input, and you must assist (Art 28). Feeding client data to AI without checking the DPA is a breach of contract before it is anything else.
Controller	You decide purposes and means — your own staff, marketing, and business records; advice you construct from client material.	The DPIA is yours, start to finish. This guide and template apply directly.
Joint / mixed	Common in practice: processor for the client's dataset, controller for your own derived records about the same people.	Split the analysis honestly in section 1 of the template. Most real engagements are mixed; pretending otherwise is where firms get caught.

THE FIVE CHECKS BEFORE CLIENT DATA TOUCHES AN AI TOOL

- **DPA check** — does the contract permit new sub-processors, and by what process? Notice period? Objection right?
- **Instruction check** — would the client recognise this as within their instructions? If you'd hesitate to describe it to them, that is your answer.
- **DPIA check** — has anyone (you or they) assessed this processing? Offer them your completed template; it converts an awkward conversation into a professional one.
- **Confidentiality check** — NDAs and professional duties (legal, accounting, health) sit on top of GDPR and are often stricter.
- **Location check** — where does the AI provider process? Your client's transfer obligations flow through you.

There is a commercial upside in this diligence: the firm arriving with a completed AI DPIA and a clean sub-processor notice wins the trust conversation against every competitor who mumbles. Client procurement teams increasingly ask for this document by name.

§ 05 / THE TEMPLATE · REPLACE EVERY [BRACKET] · EDITABLE WORD VERSION INCLUDED IN THE KIT

DPIA — [AI tool / processing activity]

AUTHOR	[name]	DATE	[YYYY-MM-DD]	STATUS	[draft / approved]
SCREENING	[triggers met — e.g. innovative technology + client personal data]			REVIEW	[+12 months]

1 · Description of the processing

THE TOOL	[name, tier, provider]
PURPOSE	[what business task, and why AI for it]
PERSONAL DATA INVOLVED	[categories, and whose — customers, staff, client's customers]
SPECIAL CATEGORY DATA?	[yes/no — if yes, which, and the Art 9 condition relied on]
OUR ROLE	[controller / processor for [client] / mixed — split described]
DATA FLOW	[in → tool → out → retained where, how long]
PROCESSING LOCATION & TRANSFERS	[regions; transfer mechanism if outside UK/EU]
SUB-PROCESSORS	[material sub-processors, from the published list]
AUTOMATED DECISIONS?	[none / advisory only / solely automated — with safeguards described]

2 · Consultation

PEOPLE AFFECTED	[views sought, or why not proportionate]
CLIENT (IF PROCESSOR)	[DPA clause; notice date; objection status]
DPO / ADVISER	[advice received, or n/a with reason]

3 · Necessity and proportionality

LAWFUL BASIS	[Art 6 basis — and Art 9 condition if special category]
WHY THIS MUCH DATA	[what is excluded and how (the three bins)]
ALTERNATIVES CONSIDERED	[manual / non-AI / anonymised — why rejected]
INDIVIDUAL RIGHTS	[how access / correction / deletion work in this tool]

4 · Risk assessment

Draw candidates from the catalogue in §03. Likelihood: rare / possible / likely. Severity (to the individual): minimal / significant / severe.

RISK TO INDIVIDUALS	LIKELIHOOD	SEVERITY	OVERALL
[e.g. client personal data absorbed into provider training]	[possible]	[significant]	[high]
[risk]			
[risk]			
[risk]			

5 · Mitigations

RISK	MEASURE	OWNER	RESIDUAL RISK
[training absorption]	[business tier · training off · setting screenshot filed]	[name]	[low]
[risk]			
[risk]			
[risk]			

6 · Outcome and sign-off

RESIDUAL RISK POSITION	[low — proceed / high — do NOT start; Art 36 consultation first]
CONDITIONS OF APPROVAL	[e.g. AMBER-bin data only; agent actions gated; review 12 months]
APPROVED BY	[name, role, date]
INTEGRATED INTO	[policy tool list (03) · supplier register · board briefing (05)]

7 · Review triggers

Re-open this DPIA — not just annually — when: provider terms or sub-processors change; tiers or features change (especially agents); data categories widen; an incident or near-miss occurs; a client's DPA changes; UK or EU rules move (as they did in 2025–26).

KEEPING IT HONEST — THREE HABITS

- **Write it before the tool goes live,** not after. A retrofitted DPIA is visible a mile away and worth half as much.
- **Let it say no sometimes.** A DPIA archive where every assessment approved everything is an archive of theatre.
- **Keep the screening records too.** The ten-minute "no DPIA needed, because…" notes are cheap and jointly they prove a working accountability process.

§ 06 / WORKED EXAMPLE

Recruitment agency, AI note-taker, client candidate calls

Condensed to show the shape and the standard of honesty — not the full form.

PROCESSING	AI transcription and summary of candidate interview calls conducted for client employers. Data: names, voices, employment history, salary expectations; occasional health disclosures (special category, incidental).
ROLE	Mixed. Processor for each client's candidate data under their DPA; controller for the agency's own candidate database records derived from the calls.
SCREENING	Triggers: innovative technology + evaluation of individuals + employment context + incidental special category data. DPIA clearly required.
KEY RISKS	Training absorption of candidate data (possible/significant); incidental health data transcribed and retained (possible/severe); summaries stating inaccuracies as fact in client reports (likely/significant); candidates unaware of AI processing (likely/significant).
MITIGATIONS	Business tier, training off, EU hosting selected; candidates told at booking and at call start, with a human-notes alternative offered; consultant reviews every summary against the recording before it reaches a client; health disclosures redacted from summaries by procedure; 30-day transcript deletion; sub-processor notice served under each client DPA with objection window.
OUTCOME	Residual risk low-to-moderate. Approved with conditions: advisory use only — no automated shortlisting; quarterly sample audit of summaries for accuracy and redaction; re-open on any provider terms change. Two clients asked for the DPIA; both accepted it as-is.

Where this fits in the kit

The **roadmap (01)** tells you when the DPIA moment arrives — the day AI touches decisions about people or client personal data. The **policy (03)** three-bin rule is your minimisation evidence for section 3. The **agents guide (04)** supplies the controls that make agent risks mitigable. The completed DPIA then feeds the **board briefing (05)** regulatory section and the **self-assessment (06)** evidence test.

WHEN TO GET HELP

This guide gets a small firm through the standard cases. Take proper advice when: special category data is core to the use rather than incidental; decisions are solely automated with significant effects; children's data is involved; residual risk stays high (Art 36 consultation territory); or an EU AI Act high-risk classification may apply to what you are deploying. An afternoon of specialist time at the right moment is cheap against any of those going wrong.

 Primary sources: UK GDPR & EU GDPR Art 35–36; ICO, *Data protection impact assessments* and *Guidance on AI and data protection*; EDPB WP248 rev.01; Data (Use and Access) Act 2025 (UK ADM reforms); EU AI Act Art 27 (FRIA). The full kit lives at peterbassill.com/ai-security-starter-kit.