NestLab

NestLab / Data Processing Agreement

Data Processing Agreement

Last updated: 2026-11-04

Download PDF

Summary

For customers who are themselves data controllers (e.g. a recruiting firm whose candidates are data subjects), we enter into a Data Processing Agreement (DPA) that complies with the EU GDPR Art. 28 and, where applicable, the UK GDPR.

This page summarises those terms and names the three large language model providers the DPA covers — Anthropic, OpenAI and Google — with the region each one processes in, the transfer mechanism, the retention term, the breach notice and the window you have to object to a new sub-processor.

What the DPA covers

- Subject matter, duration, nature, and purpose of the processing.
- Categories of data subjects and categories of personal data.
- Controller's instructions and processor's obligations.
- Confidential commitments by our personnel.
- Subprocessor list (per Subprocessors) and how to object to changes.
- Named LLM sub-processors — Anthropic, OpenAI and Google — with their processing regions, transfer mechanisms, retention terms, breach notices and objection windows.

- Technical and organizational measures (per Security).
- Data subject rights: assistance, deletion, return.
- Breach notification: 72 hours of awareness.
- Audit and inspection rights.
- International transfers: EU SCCs where applicable.
- Termination and data return / deletion.

How to get the full DPA

Email *admin [at] nestlab [dot] tech* with:

- your legal entity name and address;
- signatory name and role;
- the product(s) you intend to use.

We send a counter-signed DPA within 5 business days.

For EU AI Act Scanner customers specifically

Article 10 of the EU AI Act requires data governance for high-risk AI training data. Our DPA includes a per-system data-processing schedule that documents:

- assessment data we store (your answers, scores, generated documents);
- audit-trail data we keep (timestamps, document versions);
- retention windows (Art. 12 record-keeping);
- your obligations as data controller (Art. 10 data governance).

Data residency

Default: EU data residency for EU products — the personal data of EU customers is processed on our own infrastructure located in the European Union. Named exceptions: PropVector AI and our US-market micro-SaaS are operated from infrastructure located in Canada (North America), which is covered by the

European Commission's adequacy decision under GDPR Art. 45, so no Standard Contractual Clauses are required for that leg. Transfers to processors established outside the EU are covered by the transfer mechanism stated for that processor in the section below and in our DPA; every processor is named on the Subprocessors page.

LLM providers

Three large language model providers sit behind our products. Each of them is a sub-processor of your data under GDPR Art. 28, and each is named here with the terms its own contract commits it to. The terms below were read from the providers' own documents on 4 November 2026; the sources are listed at the end of this section.

Anthropic (Claude)

- **Where the processing happens.** Anthropic routes customer traffic to selected countries in the US, Europe, Asia and Australia by default, and stores the data in the US. Internal processing — safety review, product support, incident response — can run in any country where Anthropic or its affiliates operate. Anthropic's own cloud infrastructure providers are listed in Anthropic's sub-processor register.
- **Transfer mechanism.** The 2021 EU Standard Contractual Clauses, Module Two (controller to processor) and Module Three (processor to processor), incorporated by reference into Anthropic's Data Processing Addendum, with the UK Addendum for UK data and the Swiss Addendum for Swiss data. Anthropic's Data Processing Addendum offers no adequacy alternative, and Anthropic is not a participant in the EU-US Data Privacy Framework, so the Standard Contractual Clauses are the mechanism for this leg.
- **Retention.** Anthropic deletes API inputs and outputs from its backend within 30 days of receipt or generation. The exceptions are: a service that keeps data longer under our own control, such as the Files API; a zero data retention arrangement, which Anthropic grants per organisation, only for eligible Anthropic APIs and for Claude Code for Enterprise, and which still leaves Anthropic holding User Safety classifier results; retention needed to enforce Anthropic's Usage Policy, where a flagged conversation is kept for up to 2 years and its trust-and-safety classification scores for up to 7 years; and

retention required by law. Anthropic's Covered Models require limited data retention and review as part of its safety work.

- **Breach notice.** Anthropic notifies its customers in writing without undue delay and in any event within 48 hours of becoming aware of a security breach.
- **Objection window.** Anthropic gives its own customers 15 days from notice of a new sub-processor to object in writing; without an objection they are deemed to consent.

OpenAI (GPT)

- **Where the processing happens.** EEA and Swiss data is processed by OpenAI Ireland Ltd., and may be transferred to other OpenAI entities or to third parties outside the EEA and Switzerland in order to deliver the services. On the API, EU data residency is a project configuration rather than a default: it has to be arranged with OpenAI, it carries a 10% uplift on models released on or after 5 March 2026, it covers only the models OpenAI lists as region-supported, and it does not cover system data — account data, metadata and usage data. Requests sent to the EU endpoint terminate TLS in the selected region. Where we hold no such arrangement, content is processed in the United States and in the countries where OpenAI's affiliates and suppliers operate.
- **Transfer mechanism.** The 2021 EU Standard Contractual Clauses — Module Two where you are the controller and Module Three where you are a processor — or, where the European Commission has issued an adequacy decision under GDPR Art. 45, that decision. The UK Addendum applies to UK data. OpenAI is not a participant in the EU-US Data Privacy Framework, so the Art. 45 branch has no decision to apply here and the Standard Contractual Clauses are the mechanism for this leg.
- **Retention.** Data sent to the API is not used to train or improve OpenAI's models unless you explicitly opt in. By default OpenAI keeps abuse monitoring logs, which can contain your prompts and responses, for up to 30 days. Zero Data Retention and Modified Abuse Monitoring exist, but they are granted by OpenAI on approval and not per feature: endpoints and capabilities that OpenAI does not list as zero-data-retention eligible may still store application state, the `store` parameter is forced to false, and OpenAI may keep content in encrypted abuse monitoring logs where a model has been made ineligible for zero data retention, or for automated safety monitoring.

- **Breach notice.** OpenAI notifies its customers without undue delay after becoming aware of a personal data breach. OpenAI commits to no fixed number of hours, so we do not promise one on its behalf.
- **Objection window.** OpenAI gives its own customers 30 days from notice of a sub-processor list change to object; if the objection is not resolved within 30 days, either party may terminate the affected services.

Google (Gemini)

- **Where the processing happens.** On Paid Services Google does not use your prompts or responses to improve its products, and processes them under Google's Data Processing Addendum for Products Where Google is a Data Processor. No EU data residency is asserted for the Gemini API: the abuse monitoring log may be stored transiently or cached in any country in which Google or its agents maintain facilities. Google's terms require that only Paid Services are used where an API client is made available to users in the EEA, Switzerland or the UK.
- **Transfer mechanism.** Google LLC is an active participant in the EU-US Data Privacy Framework, so the Commission's Art. 45 adequacy decision covers that leg. Google's Data Processing Addendum names the EU-US Data Privacy Framework, the UK Extension and the Swiss-US Data Privacy Framework as transfer solutions, and the Standard Contractual Clauses as the alternative mechanism.
- **Retention.** Google logs prompts, contextual information and output for up to 55 days to detect and prevent violations of its Prohibited Use Policy. Zero data retention is not available on the Gemini Developer API by default; Google's own guidance is to use Vertex AI where guaranteed zero data retention or an enterprise data processing agreement is needed. The features that keep data regardless are: Grounding with Google Search, which stores prompts, contextual information and output for 30 days and cannot be switched off; Grounding with Google Maps, which stores the same for 30 days and cannot be switched off; the Interactions API, which stores conversation state by default unless `store` is set to false; the Live API, which retains conversation state for up to 24 hours; the File API, where files stay at rest until they are deleted or expire; and explicit context caching, which is stored for the lifetime we set. Google's implicit in-memory cache is held in RAM with a 24-hour TTL and does not count against zero data retention.

- **Breach notice.** Google notifies its partner of a data incident promptly and without undue delay. Like OpenAI, it commits to no fixed number of hours.
- **Objection window.** Google gives its partners at least 30 days' notice before a new sub-processor processes personal data, naming the sub-processor, its location and its activities, and a partner may object by written notice within 90 days of being informed.

Objection window

We notify you in writing as soon as a provider gives notice of a new sub-processor, we add it to the Subprocessors register, and you may object in writing within 10 days of our notice. Ten days is shorter than every provider's own window — Anthropic 15 days, OpenAI 30 days, Google 30 days' notice with 90 days to object — so an objection you raise reaches the provider before its own window closes.

Breach notification

We notify you without undue delay and in any event within 24 hours of becoming aware of a personal data breach. That is the chain the providers allow: Anthropic gives us 48 hours, we give you 24, and the 72 hours GDPR Art. 33 gives you to notify the supervisory authority are intact. OpenAI and Google notify without undue delay and name no fixed number of hours, so for those two the 24 hours run from the moment we are aware — we do not promise a fixed upstream deadline we cannot control.

EU AI Act Article 50

Article 50 of the AI Act has applied since 2 August 2026. It puts two duties on the provider of an AI system: to tell a person that they are interacting with an AI system, and to mark AI-generated or manipulated content in a machine-readable form. It puts a disclosure duty on the deployer for deepfakes, for emotion recognition and biometric categorisation systems, and for text published on matters of public interest without human review. NestLab is the provider of the products we ship, so the first two duties are ours and are not passed to a model vendor with the API call. Where we build a system for a customer, the customer is the deployer: we supply the disclosure wording and the marking, and the customer displays them.

Sources, read on 4 November 2026: Anthropic's Data Processing Addendum and its privacy centre articles on data retention, zero data retention and server locations; OpenAI's Data Processing Addendum and the platform documentation on data controls and data residency; Google's Gemini API Additional Terms of Service, the Gemini API documentation on abuse monitoring and zero data retention, and Google's Data Processing Addendum for Products Where Google is a Data Processor; the European Commission's guidelines on the Article 50 transparency obligations; and the US Department of Commerce's EU-US Data Privacy Framework register.

Changelog

- 2026-07-31 — Initial publication. The summary stated the EU data-residency default and that transfers to processors outside the EU are covered by EU Standard Contractual Clauses.
- 2026-08-05 — Expanded the “Data residency” section: named the EU default, the two products operated outside the EU, the GDPR Art. 45 adequacy basis for that leg, and the Standard Contractual Clauses covering transfers to processors outside the EU.
- 2026-08-15 — Removed the hosting-provider listing from the site footer. The legal pages state residency at jurisdiction level only.
- 2026-10-31 — Rewrote the “Data residency” section at jurisdiction level: no hosting provider, no data-centre location and no individual processor is named in it. The legal facts are unchanged — the EU default, the two products operated outside the EU with the Art. 45 adequacy basis, and the EU Standard Contractual Clauses for transfers outside the EU. Added a downloadable PDF of this summary.
- 2026-11-04 — Added the “LLM providers” section: Anthropic, OpenAI and Google are now named inline with the region each processes in, the transfer mechanism, the retention term and its exceptions, the breach notice, the provider's own objection window, the window we give you to object, the breach-notification chain, and the EU AI Act Art. 50 allocation. Every term was read from the provider's own document on 2026-11-04 and is sourced at the

end of the section. The “Data residency” section now points here for the transfer mechanism of each processor outside the EU.