NestLab

NestLab / Subprocessors

Subprocessors

Last updated: 2026-11-04

Download PDF

NestLab uses subprocessors to operate our products and infrastructure. This page lists the entities that may process customer data on our behalf.

This is the studio-wide register: it names every vendor that may process personal data across our products and infrastructure, and the scope stated below is the scope we operate. Each product also maintains its own detailed register with per-vendor data categories and transfer safeguards; those are available on request.

Edge & hosting

- **Cloudflare, Inc.** — CDN, edge compute, DNS, bot detection, request processing for the `/api/contact` endpoint. EU edge. DPA in place. `_cfuvid` cookie.
- **OVH SAS** — VPS hosting for Dokploy-managed app runtimes. EU products: Frankfurt (DE) and Gravelines (FR). North-America products (PropVector AI and the US-market micro-SaaS): Beauharnois, Canada — an adequacy-decision country under GDPR Art. 45. DPA in place. Also OVHcloud AI Endpoints (Gravelines, FR) for euaigap's AI risk classification - the endpoint receives the free-text description of the AI system being classified, and the description is processed on EU infrastructure; the model's answer is advisory, because the recorded classification comes from the deterministic rule engine.

Identity & authentication

- **Better Auth** — self-hosted auth library. No third-party subprocessor beyond Cloudflare's edge.
- **Supabase Auth** (CheckCV only) — JWT-based auth. DPA in place.
- **Google (OAuth / OpenID Connect)** — social sign-in and account linking for euaigap. Google returns the email address, name and OAuth profile. EU data residency is not asserted for this flow; the transfer is covered by the EU-US Data Privacy Framework adequacy decision, with the European Commission's Standard Contractual Clauses as fallback transfer mechanism.
- **Microsoft (Entra ID / OAuth 2.0 + OpenID Connect)** — enterprise single sign-on and account linking for euaigap. Microsoft returns the email address, name and OAuth profile from the identity provider's token. EU data residency is not asserted for this flow; the transfer is covered by the EU-US Data Privacy Framework adequacy decision, with the European Commission's Standard Contractual Clauses as fallback transfer mechanism.

Payments

- **Polar.sh** — Merchant of Record for all paid tiers. Stripe Connect under the hood for payouts. DPA in place. Polar acts as the data controller for payment data; we receive only the customer ID and subscription state. Polar also receives the buyer's IP address and billing country at checkout, for tax and fraud checks.

Email

- **Brevo (Sendinblue)** — three scopes. *Transactional email*: signup confirm, password reset, receipts, and contact-form delivery to `support@nestlab.tech`. *Contact records* (since 2026-10-24): a contact-form submission is also stored as a structured contact record — name, email, optional company, and the topic/product you selected — in our "NestLab - inbound leads" list, so we can track and answer enquiries instead of holding them only as email. *Subscriber lists*: the early-access waitlist — Brevo sends the confirmation and unsubscribe emails to the address you give. The subscriber list itself is held in our own database; Brevo receives no contact record for it. EU-hosted. DPA in place. See the Privacy Policy §3.4 for retention.

Analytics

- **PostHog, Inc.** — cookieless product analytics for `nestlab.tech`. EU region (`eu.i.posthog.com`, Frankfurt). PostHog processes page URLs, referrer, country, viewport size, and any custom events fired by the site. No PII, no cookies, no localStorage writes, no fingerprinting — `disable_persistence: true` + `person_profiles: "never"`. Activation is gated behind `consent.analytics === true` in the visitor's local consent record; nothing is sent before opt-in. DPA covering EU SCCs in place; data retention 90 days (PostHog default for cookieless events). See Cookie Policy §4.
- **PostHog, Inc.** — product analytics for euaigap: usage events, page views and funnel measurement. The same EU-hosted instance. It receives pseudonymous identifiers only (user, organization and AI-system UUIDs, risk tier, token counts, region, signup method, plan) — no email address and no free-text content. Client-side capture is gated behind the analytics consent category; server-side capture is not consent-gated because it carries no marketing or tracking content.

Advertising & attribution

- **Google Ads (Google LLC)** — advertising measurement and conversion attribution for euaigap. Wired but not enabled: no Ads ID is configured in any deployment, so the tag never loads and no request reaches Google. When enabled, it loads only after a marketing-consent grant; Google observes the page-load request metadata — including the client IP it sees as the requesting server, exactly as any web server does — and sets the advertising identifier, while the payloads we send carry pseudonymous identifiers only. Covered by the EU-US Data Privacy Framework adequacy decision, with the European Commission's Standard Contractual Clauses as the fallback.
- **LinkedIn (Microsoft) Conversions API** — server-side advertising conversion attribution for euaigap. Wired but not enabled: the dispatch is gated on a server-side flag and on a marketing-consent grant, and neither is configured. When enabled it receives the SHA-256 digest of the account email address, a first-party advertising identifier and the conversion event — never the email address itself, never a document and never any free-text content. Covered by the EU-US Data Privacy Framework adequacy decision, with the European Commission's Standard Contractual Clauses as the fallback.

AI providers

- **Anthropic (Claude)** — LLM via Vercel AI SDK + OpenRouter. Anthropic routes customer traffic to selected countries in the US, Europe, Asia and Australia and stores the data in the US; internal processing such as safety review, product support and incident response can run in any country where Anthropic or its affiliates operate. Transfer mechanism: the 2021 EU Standard Contractual Clauses, Modules Two and Three, incorporated by reference into Anthropic's Data Processing Addendum — Anthropic's DPA offers no adequacy alternative and Anthropic is not a participant in the EU-US Data Privacy Framework, so the Clauses are the mechanism for this leg. Retention: Anthropic deletes API inputs and outputs within 30 days of receipt or generation, with the exceptions its own terms set out — a service kept longer under our own control, a zero-data-retention arrangement granted per organisation for eligible Anthropic APIs, retention needed to enforce Anthropic's Usage Policy, and retention required by law. No fine-tuning. 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-4)** — LLM via Vercel AI SDK + OpenRouter. 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 covers the models OpenAI lists as region-supported, and it does not cover system data such as account data, metadata and usage data. 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 Standard Contractual Clauses — Module Two where the customer is the controller and Module Three where the customer is a processor — or, where the European Commission has issued an adequacy decision under GDPR Art. 45, that decision. 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 Clauses are the mechanism. Retention: data sent to the API is not used to train or improve OpenAI's models unless you opt in, and by default OpenAI keeps abuse monitoring logs, which can contain prompts and responses, for up to 30 days; Zero Data Retention and Modified Abuse Monitoring are granted by OpenAI on approval rather than per feature. No fine-tuning. 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** — LLM via Vercel AI SDK + OpenRouter. Google does not use our prompts or responses to improve its products on Paid Services, 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. Transfer mechanism: our leg is covered by the Standard Contractual Clauses OpenRouter has in place; Google LLC is an active participant in the EU-US Data Privacy Framework, and Google's own Data Processing Addendum names the Framework, the UK Extension, the Swiss-US Framework and the Standard Contractual Clauses as its transfer solutions. Retention: Google logs prompts, contextual information and output for up to 55 days to detect and prevent violations of its Prohibited Use Policy, and zero data retention is not available on the Gemini Developer API by default. No fine-tuning. Google gives its partners at least 30 days' notice before a new sub-processor processes personal data, and a partner may object by written notice within 90 days of being informed.
- **Google Gemini Developer API** — a separate, direct call from euaigap's terms-of-service check to `generativelanguage.googleapis.com`. It receives the vendor marketing pages we were asked to check and returns the vendor's own claims; it receives no personal data, no account identifiers and no uploaded documents. **No EU data residency is asserted for this flow** — for Paid Services Google logs prompts and responses for a limited period to detect and prevent violations of its Prohibited Use Policy, and that data may be stored in any country where Google has facilities. The transfer is covered by the EU-US Data Privacy Framework adequacy decision, with the European Commission's Standard Contractual Clauses as fallback transfer mechanism.
- **OpenRouter** — multi-provider LLM gateway. EU routing where available. DPA in place. The scope varies by product: our compliance product sends only the free-text AI-system description, while the chat products send the user's own conversation messages.

Scheduling

- **Cal.com, Inc.** — self-serve call booking for the `/book-a-call` page: the studio's Cal.com account holds the event type and the availability, and the booking is created there. Since 2026-11-05 the calendar opens in a window over our own page, with Cal.com's own booking interface in a frame: the visitor types their name, email address and free-text answer into Cal's own form, so those values travel from the visitor's browser straight to Cal.com and never pass through our server, and the frame's page load carries the request metadata and the IP address Cal.com logs like any web host. Cal.com sets its own cookies inside its own frame context; this site sets none for it, and no third-party script runs on this domain. Our own booking API (published for agents) is the second path: there the name, email address and process description go to our server, which passes them to Cal.com. US company; data is processed in the United States. Cal.com participates in the EU-US Data Privacy Framework (certified for non-HR data) and completes the 2021 EU Standard Contractual Clauses in its own SCC Schedule, which forms part of the Data Processing Agreement that is part of its Terms of Service. Cal.com acts as our processor for booking data and we act as its controller for its own account data — the roles are reversed from the usual. Cal.com publishes its own subprocessor list and gives notice before adding one.
- **Google Calendar and Google Meet** — the calendar the studio books studio calls into, and the video meeting the booking link opens. Paused: the studio account's meeting runs on Cal Video, and no Google Calendar is connected to it, so nothing reaches Google. The row is kept because it describes the flow that opens if the studio ever connects a studio-owned Google Calendar. While it is on: Google Calendar receives the booking details of the person who books a call — their name, email address, the meeting title, its time and its location, and the free-text message carried over from the booking panel. Google Meet carries the call itself: the participants' names and their audio and video for the duration of the meeting. Google LLC is an active participant in the EU-US Data Privacy Framework.
- **Daily, Co.** — the video platform Cal Video runs on: it carries the meeting itself, so the participants' names, their audio and video for the duration of the call, and the IP addresses of the connections that carry it. The recording and the transcription of the meeting are switched off on the studio's event type, so no recording or transcript is produced or stored. US company; data is processed in the United States. Daily, Co. is an active participant in the EU-US Data Privacy Framework (certified for non-HR data), with the 2021 EU Standard Contractual Clauses as the fallback. Daily, Co. acts as Cal.com's sub-processor

here, under the data processing agreement that forms part of the Cal.com Terms of Service; we act as controller for the booking data.

Code & CI/CD

- **GitHub, Inc.** — source code, GitHub Actions, GHCR for container images, and euaigap's immutable WORM audit-anchor repository (opt-in per tenant). Standard Contractual Clauses apply. The anchor repository is public by design — that is what makes the hash chain independently verifiable — and carries only the anchor envelope: a tenant id, an AI-system id, a Merkle root, a sequence range and a timestamp. No account identifier, no document, no free text. It is not enabled in any deployment today.
- **Cloudflare R2** — object storage for backups and static assets. EU jurisdiction option.
- **EU S3-compatible object storage** (Scaleway fr-par / OVH eu-west-par) — general-upload storage for user-uploaded files in euaigap, such as avatars and compliance evidence. EU-resident storage; no third-country transfer.

Observability

- **Beszel** — self-hosted server monitoring. No third-party.
- **GlitchTip** — self-hosted application error monitoring for euaigap. No third-party. The instance runs on our own German infrastructure (`glitchtip.euaigap.co`); the event pipeline scrubs personal data before an event leaves the host — the user email, IP address, username and id are deleted, and the request and response headers, the URL and the query string are scrubbed; an event whose sensitive pattern survives is dropped rather than sent. What remains is the error message, the stack trace and the scrubbed request metadata. No third-country transfer.
- **Uptime Kuma** — self-hosted uptime monitoring. No third-party. Triggers an internal alert on a Brevo-delivery-failure signal (when the contact form fails to deliver a message through Brevo). The alert contains only a delivery-failure signal — no message content, no email address, no name, no other personal data from the form is forwarded.

- **restic + Koofr / Backblaze B2** — encrypted backups. Self-hosted client. The backup repository holds encrypted snapshots of the production hosts, including database dumps, so it contains personal data by design; the client encrypts before upload and the repository location is EU.

Email aliases

- **Gmail (Google Workspace)** — mail reception for `admin@` / `dsar@` / `support@` / `noreply@` / catch-all. Workspace DPA in place. The `dsar@` address is the dedicated data-subject rights channel under EU GDPR Art. 15-22 (access, rectification, erasure, restriction, portability, objection); `admin@` remains the general contact and postal-address request channel.

Data-subject rights contact

EU GDPR Art. 15-22 requests (access, rectification, erasure, restriction, portability, objection) should be addressed to `dsar@nestlab.tech`. The address is delivered to the operator's Gmail inbox through Cloudflare Email Routing — no new subprocessor is added, only a forwarding rule on the same Workspace recipient already listed under "Email aliases" above.

Adding a new subprocessor

Before adding a new subprocessor that processes customer data, we:

1. evaluate the data they handle and the legal basis for transfer;
2. enter a DPA with EU SCCs if the subprocessor is outside the EEA;
3. add them to this page with a 30-day notice for substantial changes;
4. update the per-product DPA.

When a processor gives us notice of a new sub-processor of its own, we pass that notice on in writing and our customer may object in writing within 10 days of our notice. Ten days is shorter than every processor's own window — Anthropic 15 days, OpenAI 30 days, Google 30 days' notice with 90 days to object — so an objection reaches the processor before its own window closes.

The register above is our own vendor inventory. If you need the same answer for the systems you operate, that is EU AI Act compliance work — risk classification, the Annex III duty set and the documentation trail.

Changelog

- 2026-07-31 — Initial publication.
- 2026-08-01 — Widened the Cloudflare scope to request processing for `/api/contact`, and the Brevo scope to contact-form delivery. No new subprocessors.
- 2026-08-07 — Added `dsar@nestlab.tech` as the dedicated EU GDPR data-subject rights channel, delivered through Cloudflare Email Routing to the Gmail Workspace recipient already listed under "Email aliases". No new subprocessor.
- 2026-08-09 — Added the "Analytics" section with the PostHog EU entry (cookieless, consent-gated, 90-day retention).
- 2026-08-14 — Added a Cloudflare Email Routing forwarding rule for `dsar@nestlab.tech`. No new subprocessor; the same forward-to-Gmail-Workspace pattern already used for every other address in "Email aliases" above.
- 2026-09-01 — Narrowed the stated scope of the Uptime Kuma entry; removed the `sales@nestlab.tech` alias from the published alias list.
- 2026-10-24 — Widened the Brevo scope to a second scope: contact records. A contact-form submission is now stored by Brevo as a structured contact (name, email, optional company, topic, product) rather than only delivered as an email, so enquiries can be tracked and answered. No new subprocessor. In the same change, corrected the recipient in the Brevo entry from `admin@nestlab.tech` to `support@nestlab.tech`, which is where the form routes both topics.
- 2026-10-24 — Added the Google (OAuth / OpenID Connect) and Microsoft (Entra ID) entries for euaigap's account-linking flows, the Google Gemini Developer API entry for euaigap's terms-of-service check, the EU S3-compatible object storage entry for euaigap's user uploads, and widened the PostHog and Koofr entries to the scopes euaigap operates.

- 2026-10-25 — Widened three entries to the scopes our products actually operate: Polar.sh (the buyer's IP address and billing country are transmitted at checkout for tax and fraud checks), Brevo (a third scope: the early-access waitlist, whose confirmation and unsubscribe emails Brevo sends), and OpenRouter (the scope varies by product — our compliance product sends only the free-text AI-system description, while the chat products send the user's own conversation messages).
- 2026-10-26 — Added the GlitchTip entry (self-hosted error monitoring for euaigap, live) and an "Advertising & attribution" section with the Google Ads and LinkedIn Conversions API entries — both wired but not enabled in any deployment. Widened the GitHub entry with euaigap's public WORM audit-anchor repository. Driven by an internal third-party egress audit of euaigap: the product register named 11 processors while the code egresses to 15.
- 2026-10-27 — Register re-reviewed end to end against the entries above: no vendor added, removed or re-scoped since the 2026-10-26 audit. No new subprocessor.
- 2026-10-31 — Added the "Scheduling" section with two entries: Cal.com, Inc. (self-serve call booking for the booking panel on `/book-a-call`) and Google Calendar and Google Meet (the calendar a confirmed booking is written to, and the video meeting it opens). Two new subprocessors. The booking panel prefills the visitor's name, email address and free-text message into the Cal.com link, so those values now reach two US vendors; both entries state the scope and the transfer mechanism. No existing entry changed.
- 2026-11-01 — **Marked the two "Scheduling" entries as paused.** The self-serve booking panel on `/book-a-call` is switched off: the only Cal.com account available was a personal one, so a visitor of this studio site was being routed into a person's personal calendar under a personal name. With the panel off, a call request goes through our own form only and no visitor data reaches either vendor. The entries are kept (not deleted) with their scope and transfer mechanism intact, because they describe exactly what happens when the panel is switched back on — which is when the studio operates its own Cal.com account with a studio-owned calendar. No vendor added, removed or re-scoped; the register's machine-readable twin (`legal/subprocessors.json`) carries the same two rows with a `paused` flag.
- 2026-11-03 — **Booking panel live again on the studio's own Cal.com account.** The account holder is NestLab (the studio), reached at the studio's own domain address, so a visitor is no longer routed into a person's personal

calendar. The Cal.com entry is no longer paused. The meeting runs on Cal Video, which Cal.com carries through Daily, Co., so a **new entry was added** for Daily, Co. — it receives the participants' names and their audio and video for the duration of the call; the recording and the transcription of the meeting are switched off on the event type, so no recording or transcript is produced or stored. The Google Calendar and Google Meet entry stays paused: no Google Calendar is connected to the studio account, so nothing reaches Google, and the entry describes the flow that opens if one ever is. The register's machine-readable twin (`legal/subprocessors.json`) carries the same three rows and the same flags.

- 2026-11-04 — **Widened the three AI-provider entries with the terms their own contracts commit them to.** Anthropic, OpenAI and Google now state where each one processes, the transfer mechanism, the retention term and its exceptions, and the processor's own window for objecting to a new sub-processor. The entries previously said "no fine-tuning; zero data retention flag enabled where available" for Anthropic and OpenAI and "no fine-tuning" for Google, which was shorter than what the providers publish and did not name the regions the data reaches. The same change adds our own customer objection window — 10 days from our notice — to the "Adding a new subprocessor" section, and the three rows now carry the processor's own window beside it (Anthropic 15 days, OpenAI 30 days, Google 30 days' notice with 90 days to object). No new subprocessor, no vendor added or removed, and no data-flow change. The per-product DPA carries the same terms in full, with its sources.