

VERSION 2.2 · IN FORCE FROM 12 OCTOBER 2026

Data processing schedule

How we handle the personal data in your website and systems for you: only on your instructions, kept secure, with the providers we name, and returned or deleted when we finish.

In short

- You decide why and how the personal data in your website and systems is used. We handle it for you, only on your instructions, as your processor under UK data protection law.
- Your instructions are this agreement and what you approve, answer or ask for in the portal at portal.kalotect.com. We never use your personal data for purposes of our own. We use AI providers only on terms or settings that do not let them train on it, as “AI and automated work” describes, and we never train a model on it ourselves.
- We keep it secure with the measures in Annex B, which lists only what is in place today, and use only the providers in Annex C. We give you thirty days’ notice before adding a provider, and you can object. Some providers, such as image, video and voice generators and the recorder we use for meetings, get your personal data only after you approve each one.
- If something goes wrong we tell you without undue delay, and within 48 hours of finding out, and help you deal with it.
- When we finish, you choose: take your data with you, or have us delete it. We tell you in writing when the deletion is done.
- Some information we keep for ourselves, such as your team’s contact details, the record of our meetings with you and the Ledger. Our privacy notice covers that. Your customers’ personal data does not become ours because it comes up in a meeting or a message: it stays yours under this schedule.

1. What this schedule covers

This schedule forms part of the agreement between Kalotect Ltd, trading as Kalotect (“we”, “us”), and the client named in the engagement schedule (“you”). It applies whenever we handle personal data for you in providing the services, and it is the contract between a controller and a processor that Article 28 of the UK GDPR requires.

In it, “data protection law” means the UK GDPR and the Data Protection Act 2018, as amended by the Data (Use and Access) Act 2025, and any law that replaces them. “Personal data”, “processing”, “controller”, “processor”, “data subject” and “personal data breach” have their meanings in the UK GDPR. “Your personal data” means the personal data we process for you as your processor, which Annex A describes.

If the documents that make up the agreement conflict, the one higher in this list wins:

1. the special conditions in the engagement schedule;
2. this schedule, on personal data;
3. the rest of the engagement schedule;
4. our terms of business; and
5. the proposal and its plan, on how and when the work is done only.

Nothing in the engagement schedule or in our terms of business, the special conditions included, reduces the protection this schedule gives to personal data, and nothing removes or weakens a point that Article 28(3) of the UK GDPR requires. A special condition may add to that protection.

For anything about this schedule or your personal data, write to hello@kalotect.com or raise it in the portal. We have not appointed a data protection officer. Joe Kaul, head of delivery, is responsible for data protection day to day, and any departure from this schedule is decided by Michelle Wynne, director.

In plain English: This is the part of our agreement that deals with other people's information: your customers, your team and anyone else whose details are in your systems.

2. Who decides

For your personal data you are the controller and we are your processor. You decide why it is processed and on what lawful basis; we process it only to provide the services, in the way this schedule describes.

We are a controller ourselves, and not your processor, for the information we keep to run our own business, which our privacy notice describes:

- the names, roles and contact details of the people we deal with at your business, their portal accounts and sign-in records, and the signatures and initials they make in the portal, with the record we keep of each;
- proposals, the signed agreement and its certified copy, invoices, payments and the other records company and tax law require us to keep;
- the Ledger: our record of what we did for you, why, on whose instruction and with what evidence, which we keep to account for our work and to prove what was agreed. It refers to the people in your personal data by role or by reference, not by name, as "When the services end" explains;
- the record of our meetings with you: the recordings, transcripts and notes of what we discussed and agreed with your people;
- our own marketing and outreach, including the businesses and people we approach and the list of those who have asked us not to;
- security and access logs for our platform.

Your personal data does not become ours because it comes up in a meeting, a message or a comment. If your customers, patients, staff or anyone else in your personal data are named or discussed in a meeting we record, that part of the recording, transcript and notes is your personal data, which we process for you under this schedule. Please avoid naming your customers or patients in a recorded meeting. Where they are named, we remove their details from the transcript and notes we keep, and until we have, that part is returned or deleted at the end like the rest of your personal data.

We do not use your personal data for any purpose of our own, and we do not sell it or share it with anyone for their own purposes. A case study or a published result about our work for you uses your personal data only if you approve it in the portal.

In plain English: Your customers' information is yours, and we look after it for you, even when it comes up in a meeting. Information about our dealings with you is ours, and our privacy notice explains it.

3. Your instructions

We process your personal data only on your documented instructions, including about transferring it outside the UK, unless the law of the United Kingdom requires us to do something else. If it does, we tell you before we act, unless that law forbids us to.

Your documented instructions are:

- this agreement: the engagement schedule, the proposal, our terms of business and this schedule;

- what you, or anyone you have invited to act for you, approve, answer or ask for in the portal at portal.kalotect.com, each recorded in your Ledger with who gave it and when;
- any other instruction given to us in writing by someone authorised to act for you, including by email, which we confirm in the portal before we act on it.

Something one of our people records in the portal, including through support access, is your instruction only where it records an instruction you gave us in writing, with that instruction attached as evidence, or once you have confirmed it in the portal.

A question we ask you in the portal may fall to its stated default if you have not answered it by its date. A default is never your instruction to use personal data in a new way, to handle a new kind of personal data, or to send it somewhere new: for those we go ahead only on your answer, and until we have it we do not start.

We tell you straight away if we think an instruction would break data protection law, and we may pause that part of the work until it is resolved.

You are responsible for having a lawful basis for the processing you instruct, for telling the people concerned what the law requires you to tell them, for any consent you rely on, and for the accuracy of what you give us. We can draft a privacy notice or consent wording for your site as part of the services, but the decision and the responsibility stay with you.

Tell us before you give us special category data, such as health information about your patients or clients, information about criminal convictions, or information about children, so that we can agree the further measures it needs before we handle it.

In plain English: You tell us what to do, in the agreement and in the portal, and we do only that. Before your customers' information is used in a new way, we wait for a real answer from you.

4. Our people

Only people who need access to your personal data to provide the services have it, and only to the extent they need. Everyone we authorise to handle it, employee or freelancer, has committed in writing to keep it confidential, and that commitment continues after they stop working with us.

When one of us works inside your portal for you, they open support access with a stated reason. It lasts an hour at most, and your Ledger shows who came in, why, and what they did. Decisions that bind you, such as approving work, answering a question or paying, stay yours: anything one of us records through support access is your instruction only as "Your instructions" describes.

5. Keeping it secure

We take the technical and organisational measures Article 32 of the UK GDPR requires, appropriate to the risk to the people concerned. Annex B sets out the measures in place today, and the further measures we are putting in place.

We review the measures as threats and technology change. We may improve them, but we do not reduce the overall protection they give your personal data, and we tell you of any change that affects you.

6. AI and automated work

We do the work with people and AI agents, as our terms of business describe. The agents are of two kinds:

- our own agents, which run in our platform and use Anthropic's Claude models through Anthropic's API; and

- Claude Code sessions, which run in Anthropic's cloud environment, may open pull requests on your code repositories and deploy previews to Vercel, and fetch and file the notes of our meetings.

We use Claude, including Claude Code, only through accounts held by Kalotect Ltd or by our people, of two kinds. The first is on Anthropic's commercial terms and its data processing addendum, through its API or a Team or Enterprise plan; under those terms Anthropic does not train its models on what we send. The second is a Pro or Max plan, on Anthropic's consumer terms, which has a setting that lets Anthropic use what we send to improve its models: before any client work is sent to such an account, we check that the setting is off and record the check, and our platform refuses to send client work to a consumer account without that record.

The terms that apply differ between the two. On the first, Anthropic is our sub-processor under its commercial terms and data processing addendum. On the second, the contract with Anthropic is its consumer terms, made by the holder of the account, which include no data processing addendum, so Anthropic is not bound to us as a processor for that work. Our duties to you under this schedule stay as it states them in either case.

For every agent:

- each run is counted against the budget for its unit of work before it starts, and a run that would go beyond that budget waits for a person;
- an agent can never approve anything for you or for us, or move money;
- a person approves anything that carries your name, and a change reaches your live site only once Joe Kaul, or a person they have named in writing, has approved it.

For our own agents, also:

- each run has only the tools its task needs, and its tools read only the data of the client the run is for;
- what a model returns is checked against a schema before it touches your data, your files or another tool, and a run stops when it reaches its budget and asks a person;
- we screen the data a model is given for instructions hidden inside it, remove what our checks find, and log it. The screening looks for known patterns and is not a guarantee, which is why what a model returns is checked as well.

For Claude Code sessions, also:

- a session is given only the code repositories of the client its work is for;
- its work goes to a branch, a pull request and the preview built from it, and nothing it produces is merged or released until a person has reviewed and approved it;
- it reads your repositories and our meeting notes inside Anthropic's environment, where our screening for hidden instructions does not run, so we treat everything it produces as untrusted, and the report it sends back is checked against a schema before we act on it;
- a session collects new meeting recordings before they are filed to a client or checked; once a person has checked a meeting's notes, a model reads them with the names of our clients to suggest whose meeting it was, and the person who releases the plan confirms it.

When we use AI on your personal data:

- we send a model only the personal data the task needs, and use sample or test data instead wherever the work allows;
- we use AI providers only on business terms that do not allow them to train their models on what we send or what they return, or, for Claude on a Pro or Max plan, with that setting checked and recorded as off as described above, and we choose zero retention, or the shortest retention the provider offers, where it is available;

- we never use your personal data to train or fine-tune a model ourselves;
- we do not use AI to make decisions about people that have legal or similarly significant effects on them. If you ask us to build something that would, we treat it as a new instruction and help you assess it before we start;
- we send a photograph, video or recording of an identifiable person to an image, video or voice provider only when you have approved that use, with that provider, as Annex C describes.

In plain English: AI helps us work, but your customers' information is never used to train it, a model sees only what the task needs, and a person signs off anything with your name on it and anything that goes live.

7. The providers we use

You give us general authorisation to engage the sub-processors listed in Part 1 of Annex C. The providers in Part 2 receive your personal data only with your specific authorisation, given in the portal or in writing for a named use with a named provider, or, for Plaud, given as its entry in Part 2 describes. We add or replace a sub-processor only as this clause allows.

Before we add or replace a sub-processor under the general authorisation, we tell you by email, sent to the person who signed the engagement schedule for you and to every client admin, at least thirty days before the new provider receives any of your personal data, with the company that provides it, what it will do, where it will hold the data and the safeguard for any transfer. You may object on reasonable data protection grounds within that period. If you do, we will try to meet your concern, for example by not using that provider for your work. If we cannot, you may end the affected services by a message in the portal, and we refund anything you have paid for them for the time after they end.

If we have to replace a sub-processor urgently to protect your personal data or keep the services running, we replace it only with a provider of the same kind of service. We tell you before the new provider receives your personal data wherever we can, and in any case within two working days. You may object within thirty days of our notice, and if we cannot meet your concern, you may end the affected services and receive the refund described above.

We engage each sub-processor under a written contract that places on it, in substance, the data protection obligations this schedule places on us, as Article 28(4) of the UK GDPR requires, and only once we are satisfied it gives sufficient guarantees of its security. The one exception is Anthropic when we use Claude on a Pro or Max plan, which is on Anthropic's consumer terms and has no such contract, as "AI and automated work" and Annex C say. The large providers contract only on their own standard terms, which differ from this schedule in detail, for example on how quickly they tell us of a breach. Where they do, our duties to you stay as this schedule states them, and we remain responsible to you for what our sub-processors do with your personal data.

Where a service is in your own name, such as a code repository in your own GitHub account, your hosting account, your domain or your Google Analytics property, the provider is your processor under your own contract with it, not our sub-processor. We work in that account only on your instructions. We never register an account for your business, such as a domain, a social media account or a business profile, in our own name.

In plain English: We name every provider that handles your customers' information. Adding one takes a month's notice, and if you are not happy with it you can say no. Some providers are used only once you have said yes to each one.

8. Transfers outside the UK

We and our sub-processors transfer your personal data outside the UK only as Chapter V of the UK GDPR allows. You instruct us to make the transfers Annex C describes, on one of these bases:

- to a country the UK recognises as giving adequate protection, including the countries of the European Economic Area;
- to a provider in the United States that is certified under the UK Extension to the EU-US Data Privacy Framework (the UK-US data bridge);
- otherwise, under the International Data Transfer Agreement or the UK Addendum to the EU standard contractual clauses issued by the Information Commissioner, once we have assessed, and recorded, that the protection the people concerned receive is not materially lower than in the UK.

Before you sign, and again whenever we add a provider, we check which basis each transfer relies on, including whether a provider in the United States is certified on that date. If a basis stops applying, for example because a certification lapses, we move the transfer to another basis on this list or stop it, and tell you.

You can ask for a copy of the safeguard for any transfer, and of our assessment, with commercial terms removed.

9. Requests from the people concerned

If someone asks us directly to exercise their rights over your personal data, for example to see, correct or delete it, we pass the request to you within two working days and do not answer it ourselves unless you tell us to.

We help you answer requests in time to meet the deadline the law gives you, by finding, exporting, correcting, restricting or deleting the person's data in the systems we run for you, and telling you what we have done. Routine help is part of the services. Help that takes more than an hour on a single request, or that covers many requests at once, is charged at our hour rate of £120 an hour, with an estimate agreed before we start. Prices exclude VAT. We are not yet VAT registered; VAT will be added to new invoices once we are, and we will tell you before it applies.

10. If something goes wrong

We tell you without undue delay, and in any event within 48 hours, after we become aware of a personal data breach affecting your personal data, or of a security incident we reasonably believe may be one. We tell the person who signed the engagement schedule for you, every client admin and anyone else you have named to us for it, by email, and by phone where we can.

We give you as much of the following as we know, and the rest as we learn it:

- what happened and when, and when we found out;
- the kinds and approximate number of people and records concerned;
- the likely consequences;
- what we have done, and propose to do, to deal with it and reduce its effects;
- who you can talk to at our end.

We act at once to contain the breach and limit its effects, keep the evidence, and record what we did in your Ledger, describing the people affected by kind and number rather than by name. We help you decide whether to report it to the Information Commissioner's Office and whether to tell the people affected, and with what you tell them. We do not report it or tell them for you unless you ask us to or the law requires it.

In plain English: The law gives you 72 hours to report a breach to the Information Commissioner's Office unless it is unlikely to put anyone at risk. We tell you within 48, so that you have time to decide.

11. Helping you meet your other duties

Taking into account the nature of the processing and the information we have, we help you meet your duties on security, on breaches, on data protection impact assessments and on consulting the Information Commissioner's Office before high-risk processing (Articles 32 to 36 of the UK GDPR). That includes giving you information only we hold, such as how data moves through the systems we run and which sub-processors handle it.

Where the services would start processing that is likely to be high risk, such as health information at scale or a new use of AI on people's data, we tell you before we start, so that you can assess it first.

Routine help is part of the services. A larger piece of work, such as writing an assessment for you, is charged at our hour rate of £120 an hour, with an estimate agreed before we start.

12. When the services end

When we stop providing a service that involves your personal data, you choose whether we return it to you or delete it.

To have it returned, ask us in the portal or in writing at any time up to 90 days after the service ends, and we send it to you in a commonly used, machine-readable format, such as CSV or JSON, with your files in their own formats. You can also export your Ledger from the portal yourself at any time. Return in that form costs nothing; a different format, or a migration to another provider, is charged at our hour rate of £120 an hour, with an estimate agreed first.

At the end of that period, or sooner if you ask in writing, we delete your personal data and every copy we hold, and tell you in writing when it is done. Copies in our database provider's recovery history expire as that history rolls on, within ninety days at most; until then they are kept secure and used only to restore our systems.

We keep your personal data after that only where the law of the United Kingdom requires us to, for as long as it requires, and we use it for nothing else. If we have to keep any, we tell you what, under which law and for how long.

The information we keep as a controller is kept as our privacy notice says. That includes the Ledger, which we write so that it holds none of your personal data: its entries refer to your customers and the other people in your personal data by role or by reference, and evidence that holds their personal data is kept apart from it, referred to by a fingerprint of its contents, and returned or deleted with the rest. Please do not name your customers or patients in the comments and answers you give in the portal.

In plain English: Your data leaves with you or is deleted, as you choose, and we tell you in writing when it is done.

13. Showing that we comply

We make available to you the information you need to see that we are meeting this schedule. Much of it is already in your Ledger. On request we give you our current security measures, our contracts with sub-processors with commercial terms removed, and the certificates and reports we hold.

You, or an independent auditor you appoint who is bound by confidentiality, may audit or inspect our compliance with this schedule, including our systems and premises as far as they concern your personal

data:

- once a year, on thirty days' notice, at a time that does not disrupt our work for other clients;
- at any time on reasonable notice after a personal data breach affecting your personal data, or if the Information Commissioner's Office asks you to.

Before an audit we agree its scope with you, using the information above so that the audit covers what that does not answer. An audit does not give access to other clients' data or information about them. Each of us bears our own costs, unless the audit finds that we have materially failed to meet this schedule, in which case we pay its reasonable costs and put the failure right.

Audit findings are confidential between us, but nothing stops you giving them to the Information Commissioner's Office or to anyone else the law requires.

14. Liability

Each of us is responsible for meeting our own obligations under data protection law.

Our liability to each other under this schedule is subject to the limits and exclusions in our terms of business, with these differences, which apply because this schedule prevails on personal data:

- claims for breach of this schedule have their own limit, separate from the limit for other claims, as our terms of business provide;
- what our terms of business say about the loss or corruption of data does not limit or exclude a claim for breach of this schedule;
- the following are direct losses, recoverable within that separate limit: the reasonable costs of investigating a personal data breach, telling the Information Commissioner's Office and the people affected, and putting it right; and, to the extent the law allows, fines and penalties you pay to a regulator and compensation you pay to the people concerned; in each case to the extent our breach of this schedule or of data protection law caused them.

Nothing in this schedule limits either of us in a way the law does not allow, or limits what a person can claim under Article 82 of the UK GDPR. If one of us pays compensation to a person for damage that the other caused, wholly or in part, by breaking data protection law or this schedule, the other repays the share that matches its part of the responsibility.

In plain English: Claims about personal data have their own limit, separate from other claims, and the cost of dealing with a breach we caused counts as a loss we answer for within it.

15. How long this schedule lasts, and changes to it

This schedule applies from the day you sign the engagement schedule for as long as we process your personal data, including the export and deletion period after the services end. The duties on confidentiality and deletion, and the clauses on audits and liability, continue for as long as they need to.

If we break this schedule in a way that matters and do not put it right within fourteen days of your telling us, you may end the affected services at once by a message in the portal, and you keep any other rights you have.

We change this schedule only with your agreement, in the same way as the rest of the agreement, except that Annex C changes as the clause on the providers we use allows, and Annex B as the clause on security allows. If a change in data protection law means this schedule must change, we propose the change in the portal and explain why.

This schedule is governed by the law of England and Wales, and the courts of England and Wales have exclusive jurisdiction over any dispute about it.

16. Annex A: the processing

Subject matter. The services in your engagement schedule and proposal: designing, building, hosting, maintaining and improving your website and systems, and the search, social, content and care work you have asked us for.

Duration. The term of the engagement, then the export and deletion period after it.

Purpose. Only to provide those services to you, as your instructions direct.

Nature of the processing, depending on the services:

- hosting and storing your website, your systems and their data;
- building, changing, migrating and testing your website and systems, including working in your code repositories and deploying previews;
- measuring visits to your site with first-party analytics, which runs without a consent banner only where the statistics exemption in the Data (Use and Access) Act 2025 applies to your site, and asks first for anything else, such as session replay or marketing tags;
- monitoring your site's uptime and performance;
- receiving the enquiries, bookings and orders your site takes, and sending the emails your site and systems send;
- drafting, scheduling and publishing posts in your social and business accounts, and reading the replies and figures that come back;
- producing copy, images, video and audio for you, which may show or name the people you ask us to include;
- recording, transcribing and summarising our meetings with you, for the personal data about your customers, staff and others that comes up in them;
- support, fixing faults and answering tickets;
- returning and deleting the data at the end.

Categories of people, depending on the services:

- your customers, clients, patients and others who use your services, and people who enquire, book or buy through your site;
- visitors to your website;
- people who follow, message or comment on your social and business accounts, or review your business;
- your staff and contractors, and the people you ask to appear in your content;
- your suppliers, partners and other business contacts held in your systems;
- subscribers to your mailing lists.

Types of personal data, depending on the services:

- names and contact details, such as email address, phone number and postal address;
- account details for your website or systems, such as usernames and sign-in records;
- the content of enquiries, bookings, orders and messages, and order and appointment history;
- payment references, but not card details, which your payment provider handles;
- online identifiers and usage data, such as IP address, device and browser information, and the pages viewed and actions taken on your site;

- social media names, handles, profile pictures, comments and messages;
- photographs, video, voice recordings and the words of people who appear in your content;
- anything else your systems hold that the services require us to handle.

Special category data and information about criminal convictions, only where your business holds it, you have told us as the clause on your instructions requires, and we have agreed the measures it needs. Health information at a clinic or a practice is the usual example.

Your rights and obligations as controller are those this schedule and data protection law give you.

17. Annex B: security measures

These are the measures in place on the date of this schedule. The further measures we are putting in place are listed at the end; each joins this list when we tell you it is in place.

Encryption, pseudonymisation and confidentiality (Article 32(1)(a) and (b) of the UK GDPR):

- all traffic to our platform, the portal and the sites we host is encrypted in transit with TLS;
- our database, which also holds the files you add in the portal, is encrypted at rest by our database provider and held in London;
- secrets and keys are never kept in code repositories; in production they are held in our hosting provider's encrypted settings, and the tokens that start our Claude Code sessions are sealed with AES-256-GCM so that each can be opened only for the record it belongs to;
- the emails we send are stored without the tokens in their sign-in, invitation and resume links;
- people sign in to the portal with a passkey or a single-use email link that expires after fifteen minutes, and billing, settings and exports ask them to confirm again;
- we collect and keep only the personal data the services need, and where work does not need to identify anyone, such as testing, we use sample data instead.

Separation between clients:

- each client's data is kept apart in our data layer, through which every query for a client's data goes, and automated tests that try to read one client's data as another's run on every change to our platform;
- clients cannot see our internal records or other clients' data, and our own agents' tools read only the data of the client a run is for.

Integrity and accountability:

- every change we make is written to the Ledger with who made it, why and the evidence, in an append-only record whose entries are chained by SHA-256 hashes and verified every night, so that an edit cannot go unnoticed;
- a person from our studio working in your portal does so through support access, opened with a reason, limited to an hour and recorded in your Ledger;
- what our own agents' models return is validated against a schema before it is used, and every agent run is counted against a budget, as the clause on AI describes;
- every change to our platform's code runs our automated tests; dependency updates are proposed every week and merged only by a person; and a change reaches a live site only once Joe Kaul, or a person they have named in writing, has approved it.

Availability and restoration (Article 32(1)(b) and (c)):

- our database provider keeps a recovery history from which our database can be restored to an earlier point in time;

- the sites we look after are checked every five minutes, and the results are kept.

Testing and evaluation (Article 32(1)(d)):

- the tests of the separation between clients run on every change to our platform, and the Ledger's chain is verified every night, so that a failure of either is found and acted on.

People:

- a written commitment to confidentiality from everyone with access to your personal data, continuing after they stop working with us;
- access removed on the day someone stops working with us.

Measures we are putting in place, which are not yet part of the measures above. We tell you by email as each one is in place, and it then joins them:

- preview deployments and AI agent sessions that never hold credentials for production data;
- nightly backups of our own, kept for ninety days, with a restore tested every month and the result recorded;
- sign-in by passkey only for our own people, and multi-factor authentication on every provider account that holds your personal data;
- a review every quarter of who has access, removing access that is no longer needed;
- a written process for handling personal data breaches, and a test of it;
- training in handling personal data for everyone with access, before they are given it;
- continuous monitoring of our controls, and Cyber Essentials certification;
- an export of all your data from the portal in one step, and deletion at the end of the export period carried out by our systems rather than by hand.

18. Annex C: sub-processors

Part 1: sub-processors you authorise generally. For each we give the company, what it does for you and where your data is held. Where a provider in the United States holds or can reach your data, the safeguard is the UK-US data bridge if the provider is certified under it, and otherwise the UK Addendum in its data processing terms, as the clause on transfers describes. We check which applies on the day you sign, and tell you on request.

- Vercel Inc. (United States): hosts our platform, the portal, the sites we host for you and the previews of work in progress. Our platform runs in its London region; Vercel builds and serves sites through its network around the world.
- Neon, Inc. (United States): our database, with the files you add in the portal and its recovery history, held in London (Amazon Web Services, eu-west-2).
- Cloudflare, Inc. (United States): the DNS for the domains we manage for you. We do not keep your files with Cloudflare; if we start to, we tell you first as the clause on the providers we use requires.
- Resend (Plus Five Five, Inc., United States): sends email from our platform and from the sites and systems we run for you, held in the United States.
- PostHog Inc. (United States): first-party analytics for the sites we build for you, held in the European Union (Frankfurt).
- Sentry (Functional Software, Inc., United States): error reports, which can include the account and the page involved, held in the United States.
- Anthropic PBC (United States): Claude models through its API, and Claude Code sessions, which plan, draft, build and check the work and fetch and file our meeting notes, held in the United States. We use it

only through accounts held by Kalotect Ltd or by our people. Through its API or a Team or Enterprise plan, its commercial terms and data processing addendum apply. On a Pro or Max plan, its consumer terms apply, which include no data processing addendum, and we send it client work only once that plan's training setting is recorded as off, as "AI and automated work" describes.

- Google Cloud EMEA Limited (Ireland), with Google LLC (United States): our company email and documents (Google Workspace), held in Google's data centres, which may be outside the UK.
- GitHub, Inc. (United States): the code repositories in our own GitHub organisation, and the pull requests and checks on them, held in the United States.
- Zapier, Inc. (United States): carries the transcript and summary of each recorded meeting from Plaud to our platform, held in the United States.

Part 2: providers used only with your specific authorisation. They receive none of your personal data until you have approved, in the portal or in writing, a named use with that provider, or, for Plaud, until you have agreed to recording as its entry below describes. When we ask, we tell you the company that provides it, where it would hold the data, whether it trains its models on what it receives and the safeguard for any transfer, and you may say no. Work with them that involves no personal data, such as an illustration with no real person in it, does not need this approval.

- Plaud: records, transcribes and summarises our meetings with you. We record a meeting with you only once you have agreed that we may. You can agree in the portal, by answering the one-click question we put to you there; when you sign the engagement schedule, by ticking its separate, optional box; or by telling us at the start of a meeting, which the person from our studio who asked records on your record. Until you have agreed, any notes of your meetings are kept for us only, and are never planned into work or shared with you. You can withdraw your agreement in the portal at any time. Your agreement, given in any of these ways, is your specific authorisation for Plaud.
- ElevenLabs: narration, voices and transcription for audio and video we make for you.
- OpenAI, for generated images, and Google, for generated video: images and video for your content.
- Recraft, Seedream (ByteDance), Kling (Kuaishou), and the host for the Flux image models that we name when we ask: images and video for your content. We never send Seedream or Kling a photograph, video or recording of an identifiable person, so we never ask you to approve that use with them.
- Vimeo and YouTube: playing video we make for you from an account of ours, where you would rather not use your own.

Part 3: providers that are not our sub-processors. We may ask models from other providers, such as OpenAI's and Google's, for a second opinion on our own work, but we send them none of your personal data unless we have first added them to Part 1 as the clause on the providers we use allows. These are not our sub-processors for your personal data, because we use them as a controller for our own business, or because you contract with them yourself: Stripe and Starling, for our invoices and payments; Companies House and Stannp, for our own outreach; and the provider of any account in your own name, such as your domain registrar, your hosting account, a code repository in your own GitHub account or your Google Analytics property.