

VERSION 2.1 · IN FORCE FROM 12 OCTOBER 2026

How we use AI

What AI does in our work and in yours, what a person always decides, how we label what AI has made, and how you can see what it did for you.

In short

- We work with people and AI together. AI does much of the routine work: planning from meeting notes, drafting, writing and testing code, help with images and video, sorting and reporting. People direct it, check it and answer for it.
- We use the simplest tool that does the job: ordinary code first, then rules and algorithms, then small models, and large models only where they earn their place.
- A person always decides anything that carries your name, every deploy to your live site, every price and every approval. An AI agent cannot approve anything or move money, and we limit what each one can reach, as “Budgets, limits and safety” and “Claude Code sessions” explain.
- Three rules govern everything we publish. Real: nothing invented is presented as real. Disclosed: AI involvement is labelled where the law or a platform requires it. Audited: no copy goes out until it has passed our thirteen-check audit and a person has signed it off.
- We follow the transparency rules in Article 50 of the EU AI Act, which apply from 2 August 2026, the UK advertising codes and each platform’s rules on AI content. We are adding Content Credentials, under the C2PA standard, and we do not deliver an image or video we have made, or changed with AI, until it carries signed Content Credentials: a C2PA manifest, signed by us and kept with the file, inside it or alongside it.
- Our AI providers are named on this page. We use Claude, including Claude Code, only through accounts Kalotect Ltd or our people hold, either on Anthropic’s commercial terms, which do not let Anthropic train its models on what we send, or on a Pro or Max plan whose training setting we have checked and recorded as off before any client work is sent to it. We never train a model on your data ourselves.
- Our own agents work to a hard budget in pounds, and what they do for you is recorded in your Ledger. Claude Code sessions are budgeted and recorded as “Claude Code sessions” explains. If you want to know what AI did on your work, ask, and we will tell you which model was used and what it did.

1. About this page

Kalotect Ltd, trading as Kalotect, is a design and technology studio. We do our work with people and AI together. This page sets out how: which work involves AI, what a person always decides, the rules we follow on what we publish, how we label what AI has made, what happens to your data, and how you can find out what AI did on your work.

It covers our own work, the work we do for clients, and our website and portal. The promises about AI that form part of your agreement with us are in our terms of business and our data processing schedule. If this page and your agreement ever differ, your agreement is what we have agreed.

Some of what this page describes is still being built. Where it is, we say so, and we update this page when it is running.

In plain English: AI does much of the routine work, inside budgets and limits. People decide, approve and answer for it.

2. The simplest tool that does the job

We build every part of our work at the lowest rung of what we call the essential-AI ladder:

1. ordinary code, which does the same thing every time and can be tested;
2. rules and algorithms, such as a formula that scores or sorts;
3. small AI models, for narrow tasks such as sorting a brief into a few categories when rules alone cannot; and
4. large AI models, only for work that needs them, such as drafting, planning, reasoning and writing code.

Much of what we do for you involves no AI model at all. Uptime and site health checks, your Ledger and the chain that protects it, invoices, and the way we sort money in our own accounts all run as ordinary code.

Before we add a large model to any part of our work, we answer five questions in writing: what the model does that code cannot, what it costs for each unit of work, what its budget cap is, how its output is checked, and what happens when it fails. If we cannot answer them, we build it another way. We are also building a monthly review of what our AI costs against each piece of work, every run that reached its budget, and every large-model call that a simpler tool could have done.

In plain English: We use AI where it does the job better, not because it is there.

3. Where AI is part of the work

These are the kinds of work in which we use AI. In each, a person directs the work and checks what comes out of it before it reaches you or your customers.

- Meetings. We may record and transcribe our meetings with you using a Plaud recording device and its app, which also summarises them. 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. Even once you have agreed, we tell everyone present at the start, and anyone can ask us not to record, in which case we take written notes instead. Once you have agreed, the notes are filed to your record and used to plan the work. Before they are used, a person from our studio reads the notes and the transcript for sensitive details, such as anyone's health, and takes out any they find. A Claude model, in a Claude Code session, then reads them and drafts a summary for you and a plan: proposed pieces of work, questions for you and changes to your roadmap. Nothing from the meeting reaches the portal, and none of the plan becomes work, until a person from our studio has checked that the meeting is yours, read the plan and released it to you. With the release you see the summary and its actions in the portal, with the pieces of the work we show you and the plan's questions. The full transcript is for our use, and you can ask us for a copy.
- Briefs and proposals. AI sorts a new brief by the type of business and drafts a proposal from it, which a person checks, edits, prices and sends.
- Drafting. AI drafts copy, social posts, journal entries and reports in the voice you have approved. A person edits every draft, and every piece of copy passes our copy audit before it is published.
- Code. Claude Code sessions write and test code for your site and systems, as "Claude Code sessions" describes. They do not deploy to your live site.
- Images, video, voice and music. AI tools help us make and edit images, animation, video, narration and music, mostly around real photographs and footage. The rules under "What we publish: real, disclosed, audited" apply to all of it.

- Search and site care. The checks on your site are ordinary code. AI writes up what the checks find, sorts and prioritises your tickets, and proposes fixes, which a person approves.
- Reports. AI drafts your monthly report from your Ledger and your site's data, and a person checks it against the sources before it reaches you.
- Our own outreach. When we look for new clients, we check the Companies House register. A business we cannot match there as a limited company or a limited liability partnership gets a letter, a call or a visit, never an email; a small model may help match names, and when it is unsure, the answer is a letter. We are building the tools for the rest: a formula, not a model, to put businesses in order of fit, and AI to draft a first approach from a business's own website, which a person reviews before anything is sent.

4. What AI does not do

Whatever the task, an AI agent working for us does not:

- invent a quote, a review, a testimonial, a statistic, a person or a place;
- decide that something has been approved. Approval is a step in our system that only a person can take, so an agent cannot be talked out of asking;
- publish anything in your name, send anything to your customers or spend your money without your approval;
- deploy to your live site, which only Joe Kaul, or a person they have named in writing, approves;
- set a price, offer a discount or change an agreement;
- move money or change anyone's access;
- make a decision about a person that has a legal or similarly significant effect on them; or
- create a photograph of your premises, your people, your products or your work and present it as real.

In plain English: AI proposes. People decide.

5. What a person always decides

What needs a person's yes is built into our system as a fixed step and recorded in your Ledger as an approval, with who gave it and when. It is not an instruction to a model, so no model can skip it.

- You approve, in the portal at portal.kalotect.com, anything that carries your name, spends your money or reaches your customers: each page, post, advert and piece of content, design directions, objectives, plans, budget changes and your terms. Your approval covers the exact version you were shown; if it changes, we ask again.
- Joe Kaul, or a person they have named in writing, approves every deploy to a live site, on the evidence of the tests, a summary of the changes and a plan to roll back. Joe Kaul also approves every proposal before it is sent and every monthly report.
- Michelle Wynne approves our prices and discounts, any departure from our standard terms, and any exception to the content rules on this page.

When an AI agent is unsure how to proceed, it stops and asks. A question for you arrives in the portal with the options, what we recommend and a default that applies if it is not answered in time. A default never costs you more, releases a payment, publishes in your name or changes your live site.

Our own agents propose their work for a person to approve. Claude Code sessions go further: they push their work to a branch, open pull requests and deploy previews by themselves. No AI agent publishes in your name, deploys to your live site, spends money or approves anything on its own. When a person rejects an agent's work, the reason is recorded, so that the same mistake can be caught next time.

In plain English: You decide what carries your name. We decide what leaves our hands. The AI decides neither.

6. Budgets, limits and safety

Every AI agent in our platform has a profile: what it is for, the model it uses, the tools it may use and the approval it needs. Every run is counted against a unit of work, such as your project or a month of your Care, Search or Social service, and has a hard budget in pounds worked out from that unit: for most units a fixed share of its price, and for the rest a fixed sum. Before each step, the run checks that the step fits within its budget. If it would not, the run stops, records why, and asks a person, in pounds, before it goes any further. Claude Code sessions are budgeted differently, as “Claude Code sessions” explains.

The cost of the AI we use is ours and is covered by your fees. We do not pass it on to you unless your engagement schedule says so.

In our platform, we treat what a model produces as untrusted until it has been checked:

- instructions to a model and the material it works on are kept apart, and instruction-like text that our checks find in that material, for example in a web page or an email an agent reads, is removed before the model sees it and the attempt is recorded. The checks look for known patterns and are not a guarantee, which is why what a model returns is checked as well;
- anything a model returns is checked against a defined format before it can touch your data, your files or another tool, and is refused if it does not match; and
- each agent has only the tools its task needs, reads only the data of the client its run is for, and cannot raise its own permissions or move money.

In plain English: Every agent in our platform works to a budget and a set of permissions it cannot change. When it reaches either, it stops and asks.

7. Claude Code sessions

Much of our code is written in Claude Code sessions: Anthropic’s Claude, working in Anthropic’s own environment with access to code repositories, rather than inside our platform. Because they run outside our platform, some of the controls above work differently for them. This section says how.

A session works on a separate branch of your code repository, opens a pull request for a person to review, and may deploy a preview to Vercel for you to see in the portal. It does not merge its own work into your live code or deploy to your live site.

What a session can reach is set by the access we give the Claude account it runs on and the session itself, not by the separation between clients in our platform. We give a session working for you access only to your code repositories, and we check that access whenever we add a client or a repository.

A session reads your code, and our meeting notes when it is planning, inside Anthropic’s environment, where our own screening for hidden instructions does not run. So we treat what it produces as untrusted: what it reports back is checked before we act on it, and its code is reviewed by a person before it is merged.

We are building the part of our platform that starts sessions for the work planned for you. Until it is running, we start sessions ourselves. Once it is running:

- each session is given a written order for one piece of work;
- a session reports no cost of its own, so before it starts, the platform counts an allowance set for it in advance against the budget for that work, and if that would take the budget over, does not start it and asks a person. Once started, a session is not stopped part way for cost;

- the report a session sends back is checked against a defined format before we act on it;
- we record in your Ledger the order we gave the session, the report it sent back and its pull request; and
- when the platform collects new meeting recordings from Plaud, a session sees each one before it has been 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. A person can correct where a meeting is filed.

A session's own step-by-step record, such as the commands it ran, is held in the Claude account it ran on rather than in your Ledger. For any piece of work you ask about, we will look it up and tell you what it shows while it is held there.

In plain English: A session can write and test code and show you a preview. A person decides what is merged and what goes live.

8. What we publish: real, disclosed, audited

Three rules apply to everything we publish, for ourselves or for you. Work that fails one of them does not go out.

Real. Nothing invented is presented as real. If a number appears in what we write, it came from a record that names its source; if we do not have a number, we say so rather than reach for one. In particular:

- we do not create synthetic people: no face or body that has been generated or altered to look real, and no voice made to pass as a real person's; and we do not make a real person appear to say or do something they did not;
- a real, identifiable person appears only with their release on file, and a child only with a parent's consent and under the platform's stricter rules;
- we never change how your product looks: its shape, colour, texture, logo, text and proportions stay as they are. Only the setting around it may be generated, and a generated setting is labelled, as "How we label AI content" describes;
- imagery showing what a product or treatment does is real, with evidence, and never generated;
- a real place or event is shown only in real photographs or footage;
- testimonials and endorsements are real, evidenced and given with consent;
- the licence for every piece of music, stock, type and supplied material is recorded before we use it; and
- where a sector has its own rules, such as cosmetics, health, finance, alcohol or gambling, we apply them and refuse the claims they prohibit.

A synthetic voice may narrate a video as a voice-over that does not pass as any real person; where one does, the credits say so. We do not copy a real person's voice, and we do not dub a real person into words they did not speak.

Disclosed. Where the law, the advertising codes or a platform requires AI involvement to be stated, we state it on the thing itself, as "How we label AI content" describes.

Audited. No copy is published until it has passed our thirteen-check audit and a certificate exists. The checks are: truthful claims with a source; nothing invented; AI involvement disclosed where required; no unsupported superlatives or comparisons; prices with their VAT position; no dark-pattern language; plain, accessible language; British English and the right voice; correct legal names and company details; rights to every asset named; no personal data without a lawful basis; links that work and say where they go; and your own facts, such as your hours, address and prices, matching your record. Rules and code run most of the checks and a model reviews the rest. Any check the model could not complete is left for a person, and the

person who signs the certificate does so only once every check has been done. The certificate records whether a model helped write the copy and which model reviewed it.

In plain English: If it is not true, we do not publish it. If AI made it and you need to know, we say so. Nothing goes out unchecked.

9. How we label AI content

The European Union's AI Act sets transparency rules for AI in Article 50, which applies from 2 August 2026. In summary:

- people must be told when they are dealing with an AI system rather than a person, unless that is obvious;
- the providers of AI systems that generate images, audio, video or text must mark what those systems produce in a machine-readable way, so that it can be detected as artificially generated or manipulated;
- whoever uses AI to generate or manipulate image, audio or video content that resembles real people, objects, places, entities or events, and would falsely appear authentic (a "deep fake"), must disclose that it is artificial; where the content is part of an evidently artistic, creative, satirical or fictional work, the disclosure may be made in a way that does not spoil the work;
- AI-generated or manipulated text published to inform the public on matters of public interest must be disclosed, unless a person has reviewed it or exercised editorial control and someone holds editorial responsibility for publishing it; and
- the information must be clear, given at the latest when someone first sees or uses the content, and accessible.

We are a UK business, but the Act also reaches businesses outside the European Union when what their AI produces is used inside it, and much of what we make for you can be seen there. So we follow these rules in all our work, wherever it is used.

How we apply them:

- Anything generated or manipulated by AI that a viewer could take as real carries a visible label where it is seen, the platform's own AI disclosure where it has one and a line in the credits, and, once we have added Content Credentials, a machine-readable mark in them. Using such an element in a piece of work needs Michelle Wynne's approval.
- Product images on a generated setting are one such case. Your product is always the real photograph, but a generated scene or background around it could be taken as real, so each such image is labelled as having a setting made with AI where it is shown: on the image or in its caption, as each channel's rules allow, and with the platform's own AI label where it has one. Retouching that keeps a photograph true to what was in front of the camera, such as correcting colour, cleaning up marks or removing a distracting object, is not labelled. Replacing or generating the setting is.
- Work made with AI that is clearly illustrative or fantastical, such as an illustration or an animation no one would take for a photograph, is credited in a way that does not get in the way of the work.
- Text we publish on matters of public interest, such as our journal, is either written under the editorial responsibility of a named person at the studio, who reviews it in substance and is recorded as doing so, or disclosed as AI-generated. Our journal entries also say at the foot whether and how a model helped write them.
- In the portal, work done by one of our AI agents is shown under the agent's role, such as "our copywriter", with the person who reviewed it. That name alone does not make clear that it is an AI, so we are adding a plain "AI agent" label to those entries and to the questions an agent asks you. Until it is there, ask us and we will tell you whether a person or an agent did any piece of work.

- We do not run a chatbot on our public website. If we offer an AI assistant in the portal, or build one into your site, it tells people they are dealing with AI from its first message.

The UK advertising codes, written by the Committee of Advertising Practice (CAP) and the Broadcast Committee of Advertising Practice (BCAP) and applied by the Advertising Standards Authority, contain no rules specific to AI: an advert made with AI is held to the same standards as any other. An AI-generated image must not exaggerate what a product does, and a note saying AI was used does not put right an advert that misleads. For each advert we ask whether your audience would be misled if the use of AI were not disclosed, and whether a disclosure would clarify the message or contradict it.

Platforms have their own rules, and they change often. YouTube asks for realistic altered or synthetic content to be disclosed; TikTok requires AI-generated content showing realistic people or scenes to be labelled, and does not allow synthetic media of private individuals; Meta labels AI content on Instagram and Facebook. Where we publish for you, we set each platform's AI label whenever its rules require it.

We keep track of these rules and of the rules of the platforms we publish to. We are building them into a rulebook in our system, with a weekly check of the official sources, so that each piece we publish can be checked against the rules in force. When a rule we follow changes, we recheck what we have published for you in the previous twelve months and tell you if anything needs to change.

In plain English: If AI made something that could be mistaken for real, it is labelled where you see it.

10. Content Credentials

Content Credentials are an open standard, published by the Coalition for Content Provenance and Authenticity (C2PA), for a signed record, kept with an image or a video, inside the file or alongside it, that says how it was made.

We are adding Content Credentials to the images and video we make. As our terms of business commit us, we do not deliver an image or video we have made, or changed with AI, until it carries signed Content Credentials: a C2PA manifest, signed by us and kept with the file, inside it or alongside it. Until that is in place, we do not deliver one. Where an image or video on our own pages has none yet and needs an AI label, the label on the thing itself, as described above, is what tells people.

When it is in place, each image or video we make, or change with AI, will carry a record signed by us, kept with the file, that says which tool made or edited it and which parts were generated or edited with AI, and we will keep a copy of each record, with the person responsible for it, in your Ledger. Because platforms read these records to decide what to label as AI, the record has to be exact: a photograph or footage from a camera will be recorded as captured, and only the parts made with AI as generated.

Some websites and platforms strip this information when a file is uploaded, and we cannot prevent that. Where it happens, the label on the post remains. On sites we build, our site kit can show a small mark on images that carry Content Credentials.

In plain English: We are adding a signed note, kept with each file we make, inside it or alongside it, that says how it was made, so that anyone can check.

11. Your data and AI

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. Where a provider offers zero retention, or a choice of how long it keeps what we send, we choose the shortest.

We never use your data to train or fine-tune a model ourselves.

We give a model only what its task needs. Each client's data is kept apart from every other client's in our platform, and our own agents read only the data of the client a run is for. Claude Code sessions are limited as "Claude Code sessions" explains.

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, in the portal or in writing. We never send one to Seedream or Kling.

The AI providers we use, or may use for particular work, are:

- Anthropic, for its Claude models and Claude Code: our main provider, for planning, drafting, building and checking;
- Plaud, which records, transcribes and summarises our meetings;
- ElevenLabs, for narration, and for transcribing the sound in video we make for you;
- OpenAI, for generated images, and Google, for generated video; and
- Recraft, Seedream (ByteDance) and the host of the Flux image models, for images, and Kling (Kuaishou), for video.

We may later ask models from OpenAI or Google for a second opinion on our own work. We do not do so yet.

Our data processing schedule lists every provider that may handle personal data we process for you and how transfers outside the UK are protected. Before we use a provider that needs your specific approval, we tell you where it would hold the data and whether it trains its models on what it receives.

We do not make decisions about you that have legal or similarly significant effects by automated means alone. If you think an automated step has treated you or your business unfairly, write to us and a person will look at it.

In plain English: Your data is used to do your work. We never train a model on it ourselves.

12. Seeing what AI did on your work

When one of our own agents works for you, each step is recorded before and after it acts, and what it produces for you, such as a draft put to you for approval, appears in your Ledger with what was done, why, the objective it serves and the evidence. Claude Code sessions are recorded as "Claude Code sessions" explains. In the portal, each entry says who did the work: a person by name, or one of our AI agents by its role, such as "our copywriter", together with the person who reviewed it.

The portal leaves out the step-by-step record of each agent run and the model that was used. We keep them, and for any entry you ask about we will tell you which model was used and what it did.

Each piece of copy we publish has its audit certificate, which records whether a model helped write it, which model reviewed it and who signed it off.

Your Ledger cannot be quietly edited. Each entry is linked to the one before it, so any later change would show, and the chain is checked every night. You can read it in the portal at any time and export it.

If you want to know what AI did on any piece of your work, ask in the portal or write to hello@kalotect.com. We will answer from the record.

In plain English: If you ask what AI did and why, the answer is already written down.

13. AI in what we build for you

If we build an AI feature into your site or systems, such as an assistant that answers your customers' questions from content you have approved, we agree with you in writing what it may say and do, design it to tell people they are dealing with AI from its first message, and record it in your Ledger. Before we build it, we also agree in writing which of us is its provider and which its deployer under the EU AI Act, and so which of us gives the information the Act requires.

We do not add an AI feature to your site, or connect one to your customers' data, without your approval.

14. On this website

Our journal entries are drafted from the Ledger, sometimes with a model's help, edited by Joe Kaul, passed through the copy audit, and published with a line at the foot saying whether and how a model helped write them.

Case studies are written from the client's Ledger and approved by the client in the portal before they are published.

Our legal pages, this one included, are drafted with the help of AI from our own records and our code, checked by us, and reviewed by a solicitor before a director adopts them. A page that has not yet been through that review says so at the top.

15. Telling us we got it wrong

If you think something we made or published is not what it says it is, that AI was used where it should not have been, or that AI involvement was not disclosed where it should have been, write to hello@kalotect.com. We will check it against the record, correct it if it is wrong, and say on the page that we did.

16. Changes to this page

The tools, the providers and the rules all change quickly. We review this page whenever we change how we use AI, add or replace a provider, start running something this page describes as being built, or a rule we follow changes, and at least once a year. Each version has a number and a date, and we tell clients in the portal when a change affects how we work for them.