


The AI acceptable use policy you can adopt in an **afternoon**.

A ready-to-edit policy for organisations of 100 to 1,000 people: a data traffic-light, an approved tools register, rules for AI agents and a one-page staff summary, with notes on what to change and what to leave alone.



— Fig. 01 · Match the data to the tool.
Rows: green, amber and red information.
Columns: 1 public tools, 2 approved tools, 3 systems approved for one purpose. Filled: permitted.

How to use this template

Most AI policies fail in one of two ways. Legal writes them, they read like a list of prohibitions, and people quietly carry on with their personal accounts. Or they're copied from a bank, and nobody below the risk committee can follow them. We wrote this one for organisations of 100 to 1,000 people. You can adapt it in an afternoon and roll it out in a month.

Who should own it

One senior person with authority across IT, data and operations: usually the COO, CTO or CIO. Not a committee, and not Legal or IT alone. Policies owned only by Legal turn into prohibitions; policies owned only by IT turn into blocked websites. The owner needs two people in the room: whoever runs IT and security, and whoever handles data protection.

How long it takes

An afternoon for a first draft, if those three people work through it together. Allow a week or two for legal review and sign-off, and a month to roll it out properly (Appendix E).

What to customise first

- 1 Replace the placeholders. Everything in [SQUARE BRACKETS] needs a company-specific answer. Start with names and roles.
- 2 Map your data classification onto the traffic-light in section 5. If you don't have one, use ours.
- 3 Fill in the register (Appendix A) with the tools you've approved and the tools people already use. Ask Finance for AI subscriptions on expense claims; the list is usually longer than anyone expects.
- 4 Set the approval route and a target decision time in section 4. Ten working days is a sensible start.
- 5 Ask each department head for their "expensive to be wrong" list for section 7.3.

What to leave alone

Three rules do more work than the rest of the document combined: secrets and credentials never go into AI tools, the person who sends it owns it, and near misses are reported without blame. Keep all three.

Sized for you, deliberately

Enterprise frameworks are overkill for a 300-person business. You don't need an AI ethics board, a forty-page standard or a risk taxonomy. You need a short list of approved tools, a clear rule about data, a named person who answers questions within days, and a way to hear about problems early. Most staff will never read the full policy, so give them Appendix C, the policy on a page, and make the approved route faster than the workaround.

Conventions

Text in [SQUARE BRACKETS] is for you to replace. Notes marked *Guidance* are for the policy owner: use them, then delete them before you publish. Clause numbers are stable, so training material and incident reports can refer to them.

THIS IS NOT LEGAL ADVICE

This template reflects our reading of the rules as at October 2026, checked against the sources listed in Appendix G. AI and data protection law is changing quickly in the UK, the EU and the US. Ask your data protection officer or legal counsel to review your version before you adopt it, particularly clause 7.4 and section 16.

[COMPANY NAME] AI Acceptable Use Policy

POLICY OWNER	[NAME, ROLE]
APPROVED BY	[EXECUTIVE COMMITTEE OR BOARD], [DATE]
VERSION	[1.0]
EFFECTIVE FROM	[DATE]
NEXT REVIEW DUE	[DATE, NO MORE THAN SIX MONTHS AFTER THE EFFECTIVE DATE]
APPLIES TO	Everyone who works for or on behalf of [COMPANY NAME] (section 2)
RELATED POLICIES	[INFORMATION SECURITY POLICY], [DATA PROTECTION POLICY], [DATA CLASSIFICATION POLICY], [IT ACCEPTABLE USE POLICY], [PROCUREMENT POLICY], [DISCIPLINARY POLICY]
QUESTIONS	[AI POLICY CONTACT, FOR EXAMPLE AN EMAIL ADDRESS OR CHAT CHANNEL]

GUIDANCE FOR THE POLICY OWNER • DELETE BEFORE PUBLISHING

Set the review at six months, not twelve. The tools, their terms and the rules around them change faster than an annual cycle.

Contents

01	Purpose and principles	3	13	Incidents and near misses	13
02	Scope	3	14	Training and acknowledgement	14
03	Roles and responsibilities	4	15	Breaches of this policy	14
04	Approved tools	5	16	Legal and regulatory context	14
05	Data: what may go into which tools	6	17	Review and version control	15
06	Acceptable and prohibited uses	7	Appendix A	Approved tools register	16
07	Human review and accountability	8	Appendix B	Data traffic-light quick reference	18
08	AI agents and automation	9	Appendix C	The policy on a page	19
09	Transparency and disclosure	10	Appendix D	Tool request form	20
10	Intellectual property and confidentiality	11	Appendix E	Rollout checklist for the policy owner	22
11	Security	12	Appendix F	Glossary	24
12	Cost and accounts	12	Appendix G	References	25

1. Purpose and principles

- 1.1 [COMPANY NAME] wants everyone to use AI tools to do better work, faster. This policy explains how to do that without putting our customers, colleagues, information or reputation at risk.
- 1.2 Five principles sit behind every rule in this policy:
 - (a) **Use approved tools for company work.** If the tool you need isn't approved, ask for it (section 4). Don't work around the process.
 - (b) **Match the data to the tool.** The data traffic-light in section 5 decides what information may go into which tools.
 - (c) **A human owns every output.** If you use it, send it or act on it, you are accountable for it.
 - (d) **Be open about AI use** where it matters to the people affected.
 - (e) **Report problems early.** Nobody is blamed for reporting a near miss.
- 1.3 This policy sits alongside our [INFORMATION SECURITY POLICY], [DATA PROTECTION POLICY] and [IT ACCEPTABLE USE POLICY]. Where they overlap, follow the stricter rule.
- 1.4 This policy does not form part of any contract of employment or engagement, and [COMPANY NAME] may change it at any time.

GUIDANCE FOR THE POLICY OWNER • DELETE BEFORE PUBLISHING

Keep the principles short enough to remember. Most staff will only read Appendix C, so the principles there must match these word for word.

2. Scope

- 2.1 This policy applies to all employees, contractors, consultants, temporary and agency workers, interns, and anyone else who works for or on behalf of [COMPANY NAME] or has access to our systems or information ("you").
- 2.2 An "AI tool" is any software that uses machine learning to generate, summarise, transform, classify or act on content. It includes:
 - (a) general-purpose assistants and chatbots, such as [LIST THE ONES YOUR STAFF USE];
 - (b) AI features built into software we already use, such as meeting summaries, writing assistants in email and documents, and AI features in our [CRM, HR AND FINANCE] systems;
 - (c) coding assistants and command-line coding agents;
 - (d) AI agents: systems that take actions, such as sending messages, updating records, running code, browsing the web or making purchases (section 8);
 - (e) browser extensions, plug-ins, connectors, and desktop or mobile apps with AI features, including notetakers that join calls;
 - (f) models used directly through an API, and open-weight models run on company or personal hardware;
 - (g) tools that generate or alter images, audio, voice or video.

- 2.3 Switching on a new AI feature in software we already use counts as adopting a new AI tool.
- 2.4 This policy applies whichever account you use. Personal accounts, including free tiers and personal paid subscriptions, may only be used for company work with Green information (section 5).
- 2.5 It applies on company and personal devices, wherever you are working.
- 2.6 This policy covers how we use AI. If [COMPANY NAME] builds AI into products or services it sells, that work must also follow [AI DEVELOPMENT STANDARD OR PRODUCT RISK PROCESS] .

GUIDANCE FOR THE POLICY OWNER • DELETE BEFORE PUBLISHING

Clause 2.3 is easy to overlook, and it matters. AI arrives through software updates far more often than through new purchases, and a transcription feature switched on for every meeting is a bigger data decision than one person's chatbot account. If you sell AI-enabled products, you need more than this policy: evaluation, testing and, if you sell into the EU, possibly obligations as a "provider" under the EU AI Act. That belongs in a separate document.

3. Roles and responsibilities

- 3.1 These roles are responsible for this policy:

ROLE	WHO	RESPONSIBILITIES
AI Policy Owner	[NAME, ROLE, FOR EXAMPLE COO OR CTO]	Owns this policy and its reviews. Settles questions the policy doesn't answer. Reports on AI use, incidents and spend to [EXECUTIVE COMMITTEE] every [QUARTER] .
Tool Approvers	[HEAD OF IT OR SECURITY] and [DATA PROTECTION LEAD] , with [LEGAL] for contract terms	Assess requests for new tools, features and connections (section 4). Decide the highest data class each tool may handle.
AI Register Keeper	[NAME, ROLE, FOR EXAMPLE IT OPERATIONS MANAGER]	Keeps the registers in Appendix A accurate and published. Logs incidents and near misses.
Data Protection Lead	[DPO OR NAMED LEAD]	Advises on personal data, data protection impact assessments (DPIAs), automated decisions and transparency.
Budget Owners	[FINANCE DIRECTOR AND DEPARTMENT HEADS]	Approve paid AI accounts, subscriptions and API spend (section 12).
Line managers	All people managers	Make sure their teams know this policy and complete the training. Define their area's "expensive to be wrong" list (section 7.3). Encourage near-miss reporting.
Everyone	All workers	Follow this policy, complete the training, ask when unsure, and report incidents and near misses.

- 3.2 One person may hold more than one role, but nobody should be the only approver of a tool they requested.

In a 300-person business, three or four people usually cover all of these roles. A pattern that works: the COO or CTO owns the policy, the head of IT keeps the register, and whoever handles data protection approves anything touching personal data. Only set up an AI steering group if you already run steering groups that make decisions quickly.

4. Approved tools

- 4.1 For any information above Green (section 5), use only tools listed on the Approved Tools Register (Appendix A), and only for the uses and data classes listed against each tool. The current register is at [\[LINK\]](#).
- 4.2 To request a new tool, a new AI feature in an approved tool, or a new connection between an AI tool and another system, complete the Tool Request Form (Appendix D) and send it to [\[EMAIL ADDRESS OR TICKET QUEUE\]](#). The approvers aim to decide within [\[10\]](#) working days.
- 4.3 The approvers assess each request against these questions. A “no” or “don’t know” doesn’t automatically mean refusal; it may limit the data class or the uses the tool is approved for.
- (a) **Training use.** Does the provider use our inputs or outputs to train or improve its models? Is that excluded by contract, and not just by a setting each user can change?
 - (b) **Retention.** How long are prompts, files and outputs kept? Can we shorten that, or choose zero retention?
 - (c) **Region.** Where is data processed and stored? If personal data leaves the UK [\[OR EEA\]](#), what transfer safeguard applies?
 - (d) **Terms.** Are business or enterprise terms in place, including a data processing agreement where personal data is involved, confidentiality, and ownership of outputs?
 - (e) **Identity.** Does the tool support single sign-on (SSO) and central account management, so access ends when someone leaves?
 - (f) **Logging.** Can administrators see usage and audit logs, and export them to [\[LOGGING OR SECURITY TOOL\]](#)?
 - (g) **Sub-processors.** Which other companies, including model providers, process our data? Will we be told when that list changes?
 - (h) **Admin controls.** Can we control features centrally, such as connectors, plug-ins, web access, file upload, sharing and agent actions?
 - (i) **Security assurance.** Does the provider hold relevant certifications or reports, such as ISO/IEC 27001 or SOC 2 Type II?
 - (j) **Exit.** Can we export our data, and have it deleted when we leave?
 - (k) **Cost.** Is pricing per seat, based on usage, or both? Can we set spending limits and alerts?
- 4.4 Each decision is recorded on the register as Approved (with the highest data class and permitted uses), Approved with conditions, Under review, or Not approved. A refusal names an approved alternative wherever one exists.

- 4.5 Approvals are reviewed at least [ANNUALLY], and whenever a provider materially changes its terms, ownership, sub-processors or features.
- 4.6 Trials and pilots follow the same rules. Until a tool is approved, use it with Green information only.

GUIDANCE FOR THE POLICY OWNER • DELETE BEFORE PUBLISHING

Make the approved route faster than the shadow one. If approval takes two months, people will use personal accounts and you'll never hear about it. Eleven questions are enough for a mid-market business. Many providers answer most of them in a public trust centre or their standard data processing agreement, so a low-risk tool can often be cleared in a day. Save the long vendor questionnaire for tools that will touch Red information or take actions.

5. Data: what may go into which tools

- 5.1 Before you put information into an AI tool, whether by typing, pasting, uploading a file, connecting a data source or letting an agent read it, check its classification under our [DATA CLASSIFICATION POLICY] and apply the traffic-light below.

LIGHT	CLASSIFICATION	EXAMPLES	WHERE IT MAY GO
● Green	[PUBLIC]	Our website and published marketing; public filings; job adverts; general questions that contain no company or personal details	Any AI tool, including personal accounts. Approved tools are still preferred.
● Amber	[INTERNAL] and [CONFIDENTIAL]	Internal emails, documents and meeting notes; plans and non-public financials; source code; contracts; customer and supplier details; personal data about customers or staff, other than special category data	Approved tools only (Appendix A), on a company account, within the data class the register allows. Personal data only in tools with a data processing agreement in place.
● Red	[RESTRICTED]	Passwords, API keys, tokens and private keys; payment card data; special category personal data (such as health, ethnicity, religion, sexual orientation, biometric and genetic data) and criminal records data; legally privileged material; inside information and unannounced deals; details of security vulnerabilities; anything a client contract forbids us to share with third parties or AI services	No AI tool, except a system the register lists as approved for that specific information and purpose. Secrets, credentials and payment card data: never, with no exceptions.

- 5.2 If you're unsure which light applies, treat the information as the higher class and ask [AI POLICY CONTACT].
- 5.3 Share the minimum. Remove names, account numbers and other details the tool doesn't need. Removing names doesn't always make data anonymous, so the rules for personal data may still apply.
- 5.4 Information you receive from customers, suppliers or partners gets the class it would have if it were ours. Contracts may restrict its use further (section 10).
- 5.5 Connecting an AI tool or agent to a system, such as a mailbox, shared drive, CRM or code repository, gives it access to everything your account can see there, not just the item you have in mind. Classify

the connection by the most sensitive information the tool could reach.

GUIDANCE FOR THE POLICY OWNER • DELETE BEFORE PUBLISHING

If you already have a classification scheme, map its levels onto the three lights and replace the bracketed names. If you don't, these four levels are a sound place to start. Clause 5.5 matters more every month: a connected assistant inherits the user's permissions, and few organisations have tidied those recently. The most valuable line in this policy is "secrets and credentials: never". One AI gateway provider has described scanning the traffic through its own service and finding names, email addresses and live access tokens on their way to model providers.

6. Acceptable and prohibited uses

6.1 Within the rules in sections 4 and 5, we encourage you to use AI tools for:

- drafting and editing emails, documents, presentations and marketing copy, which you then review;
- summarising long documents, meeting notes and research, and checking the summary against the source for anything you will rely on;
- brainstorming, outlining and getting a first draft started;
- analysing data and building spreadsheets and formulas, with the numbers checked;
- writing, explaining, reviewing and testing code in approved coding tools;
- translating, with a fluent human check for anything customer-facing or contractual;
- searching our own knowledge through approved tools;
- learning a new subject or preparing for a meeting.

6.2 These uses are allowed but need more care. Each needs human review under clause 7.3, and some need specific approval:

- (a) **Customer communications** and anything published externally.
- (b) **Recruitment and people management.** AI may help draft job adverts, role descriptions and interview questions. It must not be used to screen, rank, score or reject candidates, or to inform decisions about employees' performance, pay, promotion, discipline or dismissal, unless that specific use has been approved by [AI POLICY OWNER] and [DATA PROTECTION LEAD] after a data protection impact assessment.
- (c) **Decisions about customers' eligibility,** including credit, pricing, insurance or access to services. The same approval applies.
- (d) **Factual, legal, financial, medical or safety claims,** which someone competent to judge them must check against a reliable source.
- (e) **Code that will run in production,** which goes through normal code review and testing.

6.3 You must not use AI tools to:

- (a) put Amber information into a tool not approved for it, or Red information into anything other than a system approved for that specific use;
- (b) make, or fully automate, decisions with legal or similarly significant effects on people, except through a process approved under clause 7.4;

- (c) give legal, tax, financial, medical or other regulated advice to anyone without review by a suitably qualified person, or treat AI output as legal advice to [COMPANY NAME] ;
- (d) create images, video, audio or voice depicting real, identifiable people (including colleagues, customers and public figures) without their consent, or create content intended to deceive. Sexual or intimate imagery of any real person is prohibited in all circumstances;
- (e) impersonate a real person, or let an AI system present itself as a human (section 9);
- (f) infer the emotions of colleagues or job candidates, for example by analysing their faces, voices or messages;
- (g) monitor or profile colleagues outside processes approved by [HR] and [DATA PROTECTION LEAD] ;
- (h) produce discriminatory, harassing, defamatory or otherwise unlawful content, or anything that breaches our [CODE OF CONDUCT] ;
- (i) get around security controls, tool safeguards or this policy, including by using a personal account to avoid a block;
- (j) record or transcribe calls and meetings without telling the participants (clause 9.4);
- (k) run an agent or automation that acts on our systems, or contacts anyone outside [COMPANY NAME] , without registering it first (section 8).

6.4 If you're unsure whether a use is acceptable, ask [AI POLICY CONTACT] before you go ahead.

GUIDANCE FOR THE POLICY OWNER • DELETE BEFORE PUBLISHING

Resist the urge to ban everything. A blanket ban pushes usage onto personal phones, where you can't see it at all. This list prohibits what can do real damage and leaves the rest to the traffic-light. Recruitment, people management and eligibility decisions are where AI use is most heavily regulated. The EU AI Act treats AI used for recruitment, managing workers, credit scoring, and life and health insurance pricing as high-risk, with obligations due from 2 December 2027 on the current timetable. These decisions are also subject to the UK's automated decision-making rules, and in the US laws such as New York City's Local Law 144 cover hiring tools. Clause 6.3(f) mirrors the EU AI Act's ban on emotion recognition in the workplace, which has applied since 2 February 2025 with narrow medical and safety exceptions; we suggest applying it everywhere. Take advice before approving any use under 6.2(b) or 6.2(c).

7. Human review and accountability

- 7.1 The person who uses, sends, publishes or acts on AI output is accountable for it, exactly as if they had produced it themselves. If you send it, you own it. [COMPANY NAME] will not accept "the AI wrote it" as an explanation.
- 7.2 AI tools make mistakes, including confident and plausible errors, invented references, out-of-date information and faulty arithmetic. Check facts, figures, names, quotations, citations and code before you rely on them.
- 7.3 Where being wrong is expensive, a competent person must review AI output before it is used. As a minimum, this covers:
 - (a) anything sent to customers, regulators, investors or the media, or published;
 - (b) contracts, legal and regulatory documents, financial reporting and board papers;

- (c) code deployed to production, infrastructure changes and security configuration;
- (d) anything affecting health, safety or people's rights;
- (e) [ADD THE AREAS SPECIFIC TO YOUR BUSINESS] .

A review means reading for substance: the reviewer should be able to explain and defend the output.

- 7.4 [COMPANY NAME] does not make decisions about individuals based solely on automated processing that have legal or similarly significant effects on them, unless the specific process has been approved by [AI POLICY OWNER] and [DATA PROTECTION LEAD] after a data protection impact assessment, and includes the safeguards the law requires. Under the UK GDPR (Articles 22A to 22D, in force since 5 February 2026), those safeguards include telling people about such decisions, letting them make representations, giving them human intervention and letting them contest the decision. Stricter rules apply where special category data is involved, and the EU GDPR may also apply to decisions about people in the EU.
- 7.5 Human involvement in a decision must be meaningful. The reviewer needs the authority, information and time to change the outcome. Approving the system's output by default does not count.
- 7.6 Where AI materially informed a significant business decision, record which tool was used and how its output was checked [IN THE DECISION PAPER OR PROJECT FILE] .

GUIDANCE FOR THE POLICY OWNER • DELETE BEFORE PUBLISHING

Ask each department head for their "expensive to be wrong" list in a single meeting. It's a better use of an hour than a risk matrix. Clause 7.4 summarises the law; it doesn't replace advice. The UK rules changed in February 2026, and when we checked in October 2026 the ICO's updated guidance on automated decision-making was still being finalised. Check the current position with your data protection officer or legal counsel before approving any automated decision about people.

8. AI agents and automation

- 8.1 An AI agent is an AI system that takes actions as well as producing content: for example sending messages, changing records, running code, moving money, filling in web forms, or calling other systems through connectors and APIs. This section also covers scheduled automations that use AI.
- 8.2 Every agent or AI automation that acts on our systems or contacts anyone outside [COMPANY NAME] must be entered on the agents register (Appendix A) before it runs, with a named human owner, what it can read, what it can do, which credentials it uses, its spending limit and how to stop it.
- 8.3 Give agents the least access they need. Use their own scoped, short-lived credentials rather than a person's account or long-lived keys. Start read-only and add permissions only when needed. Agents under development must not touch production systems or live customer data.
- 8.4 A human must approve these actions before an agent carries them out:
- (a) sending external communications, other than approved templates;
 - (b) payments, refunds, purchases, credit or discounts;
 - (c) deleting or bulk-changing records;
 - (d) changes to production systems, access rights or security settings;

(e) anything listed in clause 7.3;

(f) [ADD OTHERS] .

- 8.5 Never put passwords, API keys or tokens in prompts, agent instructions, or configuration files kept in code repositories or shared documents. Store them in [SECRETS MANAGER] . If a credential is exposed to an AI tool, rotate it straight away and report it (section 13).
- 8.6 Treat content from outside [COMPANY NAME] as untrusted. Web pages, emails, documents, tickets and the output of other tools may contain hidden instructions (prompt injection). Do not give a single agent all three of the following unless [HEAD OF IT OR SECURITY] has approved the controls: access to sensitive information, exposure to untrusted content, and the ability to send information outside the organisation.
- 8.7 Agents must log their actions, the tools they called and the data they accessed, somewhere the owner and [IT OR SECURITY] can review. Keep logs for [90 DAYS OR AS SET BY OUR RETENTION POLICY] .
- 8.8 Every agent must have a documented way to stop it and revoke its access quickly, known to its owner and to [IT] .
- 8.9 Where a control in this section can be enforced by technology, such as permissions, spending caps, network allowlists or approval steps, enforce it in the tooling. Don't rely on written rules alone, and don't describe a control in this policy, in agent instructions or in documentation unless it actually operates.
- 8.10 Coding assistants and coding agents may work in [APPROVED REPOSITORIES] . They must not hold production credentials. Modes that let an agent run commands without asking are allowed only in isolated environments [SUCH AS A CONTAINER OR VIRTUAL MACHINE WITH NO PRODUCTION ACCESS] . All AI-written code goes through normal review.

GUIDANCE FOR THE POLICY OWNER • DELETE BEFORE PUBLISHING

This is the section that will matter most over the next year. Clause 8.9 addresses a real failure mode: agents act on the controls they are told about, not the ones that exist. In the July 2026 incident in which OpenAI's evaluation agents broke into Hugging Face, researchers who studied the transcripts found the agents spent days trying to evade a reviewer that the published documentation described but the real scorer never used. If your policy says something is monitored, make sure it is. Clause 8.6 names the combination that turns prompt injection from an embarrassment into a breach.

9. Transparency and disclosure

- 9.1 Tell people when they are dealing with an AI system rather than a person, for example a chatbot or an AI voice agent, unless that is obvious from the context. Never let an AI system claim to be a human or a named member of staff.
- 9.2 Label images, audio and video that AI has generated or materially altered when they show real people, places or events and could be mistaken for genuine. If we publish AI-generated text to inform the public on matters of public interest, disclose that AI was used, unless a person has reviewed it and [COMPANY NAME] takes editorial responsibility for it.

- 9.3 You don't need to label routine AI assistance, such as an email you drafted with AI, then reviewed and sent under your own name. Clause 7.1 already makes it yours.
- 9.4 Announce AI notetakers at the start of a meeting. For meetings with external participants, get their agreement first, and don't record if anyone objects. Share AI summaries only with people entitled to see the meeting content.
- 9.5 Be open with colleagues about AI use where it affects how they will rely on your work, for example "first draft generated with [TOOL], checked against the signed contract".
- 9.6 [DATA PROTECTION LEAD] makes sure our privacy notices describe how we use personal data in AI tools where the law requires it.
- 9.7 If a customer contract requires us to disclose AI use, or restricts it, the contract takes precedence.

GUIDANCE FOR THE POLICY OWNER • DELETE BEFORE PUBLISHING

The EU AI Act's transparency rules (Article 50) have applied since 2 August 2026. Most of them fall on the providers that build AI systems, but organisations using AI have duties too, notably disclosing deepfakes and AI-generated text published to inform the public. Clauses 9.1 and 9.2 follow that model, and they're sensible practice even with no EU exposure. Clause 9.3 stops the policy turning into a labelling exercise nobody follows.

10. Intellectual property and confidentiality

- 10.1 Put client or third-party information into an AI tool only where our contract with them allows it and section 5 permits. Some client contracts prohibit the use of AI or of new sub-processors. [LEGAL OR ACCOUNT OWNERS] keep a list of these clients at [LINK].
- 10.2 Don't upload material we don't have the right to use in this way, such as paid research reports, licensed datasets or other organisations' confidential documents, where the licence or agreement forbids it.
- 10.3 Work you produce with AI tools in the course of your job belongs to [COMPANY NAME] [AS SET OUT IN YOUR CONTRACT]. Copyright protection for AI-generated material is uncertain and differs between the UK, the EU and the US. Where ownership matters, for example brand assets, product code or published content, make a substantial human contribution and keep a record of it.
- 10.4 Don't ask AI tools to reproduce copyrighted works, logos or trade marks, or to imitate a named living artist or writer for commercial use. Check images for recognisable third-party marks or likenesses before publishing them.
- 10.5 AI-written code goes through the same review, licence checks [USING TOOL] and security scanning as any other code. Where an approved coding tool can block suggestions that match public code, switch that setting on.
- 10.6 Confidential information stays confidential whether or not an AI tool was involved in handling it.

GUIDANCE FOR THE POLICY OWNER • DELETE BEFORE PUBLISHING

Ask Legal for the list in clause 10.1 before you launch. Staff can't follow a contractual restriction they've never heard of.

11. Security

- 11.1 Use approved AI tools through single sign-on on [MANAGED DEVICES] . Don't use Amber information with AI tools on unmanaged personal devices [UNLESS OUR BRING YOUR OWN DEVICE POLICY ALLOWS IT] .
- 11.2 Never put passwords, API keys, tokens or other secrets into prompts or into files you share with an AI tool (clause 8.5).
- 11.3 Content from outside can carry hidden instructions that make an AI tool leak information or take unwanted actions (prompt injection). Take care when you ask an AI tool that can see company information to process emails, web pages or documents from outside. Be suspicious of AI output that asks you to click a link, share information or change a setting.
- 11.4 Don't install AI browser extensions, plug-ins, connectors (including Model Context Protocol servers), desktop apps or mobile apps for work unless they are on the register. Many can read everything on your screen or in the connected system.
- 11.5 Don't create public or shareable links to AI conversations that contain Amber or Red information. Anyone with the link may be able to see them.
- 11.6 Don't download and run AI models, datasets or notebooks from public sources on company systems without approval from [IT OR SECURITY] . Some file formats can run code when they are loaded.
- 11.7 AI makes convincing phishing emails, cloned voices and fake video cheap to produce. Verify unusual requests, such as payments, changes to bank details, password resets or data transfers, through a separate channel you already trust, even if the voice or face seems familiar.
- 11.8 [IT] may block or monitor access to AI services that are not approved, in line with our [IT ACCEPTABLE USE POLICY] and [MONITORING NOTICE] .

GUIDANCE FOR THE POLICY OWNER • DELETE BEFORE PUBLISHING

Clause 11.6 is worth a conversation with your engineering and data teams: the July 2026 intrusion into Hugging Face began with a malicious dataset uploaded for processing. On 11.8, block unapproved AI services only once an approved alternative exists. Blocking first and offering nothing pushes people onto their phones.

12. Cost and accounts

- 12.1 Only [ROLES, FOR EXAMPLE IT AND FINANCE] may create paid AI accounts, subscriptions or API keys for company use. Don't pay for AI tools for company work on a personal or company card unless the tool is approved and the spend has a budget owner.
- 12.2 Every paid AI tool and API key has a named budget owner, recorded on the register.
- 12.3 Tools charged by usage, including APIs, agents and some coding assistants, must have spending alerts set at [50%, 80% AND 100%] of budget. Unattended agents and automations must also have a hard spending limit.

- 12.4 API keys are issued per team or application, never shared between people, stored in [SECRETS MANAGER] and revoked when no longer needed.
- 12.5 [FINANCE] reviews AI spend by team and tool every [MONTH] against the outcomes it was meant to deliver.

GUIDANCE FOR THE POLICY OWNER • DELETE BEFORE PUBLISHING

Pricing for many AI tools now depends on usage rather than seats, and an unattended agent can run up a bill overnight. Measure what the spend bought, such as tickets resolved or pull requests merged. Token counts say nothing about value, and ranking staff by AI usage rewards people for burning tokens.

13. Incidents and near misses

- 13.1 Report immediately to [CHANNEL, EMAIL ADDRESS OR SERVICE DESK] if:
- (a) Amber or Red information has gone into a tool not approved for it, even by mistake;
 - (b) a password, key or token has appeared in a prompt, an output, a log or a shared conversation;
 - (c) an AI tool or agent has taken, or tried to take, an action nobody intended;
 - (d) incorrect AI output has reached a customer, a regulator or the public, or informed a significant decision;
 - (e) you suspect prompt injection, a malicious extension or a compromised account;
 - (f) an AI tool has produced discriminatory, offensive or unlawful content in a work context.
- 13.2 Report near misses too: something that could have gone wrong but was caught in time, such as a wrong clause spotted at review, an agent stopped before it sent an email, or a file nearly uploaded to the wrong tool. Use the short form at [LINK].
- 13.3 Promptly reporting a near miss or an honest mistake will not in itself lead to disciplinary action, and we will take it into account in your favour. Deliberate or reckless breaches, and attempts to hide an incident, are treated differently (section 15).
- 13.4 If an incident involves personal data, [DATA PROTECTION LEAD] decides whether it must be reported to the Information Commission's Office (ICO). For a notifiable breach, that means within 72 hours of becoming aware of it, where feasible.
- 13.5 The AI Register Keeper logs all incidents and near misses. The AI Policy Owner reviews them every [QUARTER] and changes the policy, the register or the controls in response. Each review looks for patterns and for controls that need to change.

GUIDANCE FOR THE POLICY OWNER • DELETE BEFORE PUBLISHING

Aviation's near-miss reporting works because people aren't punished for reporting. Without a channel for AI near misses, you tend to find out about problems from customers. A shared inbox and a two-minute form are enough to start. Count the reports. A silent channel in an organisation that uses AI every day usually means people don't trust it.

14. Training and acknowledgement

- 14.1 Everyone completes [AI BASICS TRAINING] within [30 DAYS] of this policy taking effect or of joining, and a refresher every [YEAR].
- 14.2 Training covers this policy, the data traffic-light, how AI tools fail (errors, bias and prompt injection), how to review output, and how to report incidents and near misses.
- 14.3 People in higher-risk roles get extra training for their role. That includes anyone who builds or runs agents, and people in HR, finance, legal and customer-facing teams who use AI regularly.
- 14.4 Everyone confirms they have read and understood this policy [THROUGH THE HR SYSTEM] when they join and after each significant update. [HR] keeps records of training and acknowledgements.

GUIDANCE FOR THE POLICY OWNER • DELETE BEFORE PUBLISHING

If you have staff in the EU, or your AI outputs are used there, Article 4 of the EU AI Act expects you to take measures to support your staff's AI literacy. Amendments in force since 27 July 2026 softened the wording but kept the duty on organisations using AI. Recorded training against this policy is a practical way to support that duty and to keep evidence of what you've done. Keep the training short and practical: thirty minutes on real examples from your own business will do more than a two-hour module on how large language models work.

15. Breaches of this policy

- 15.1 Breaches may lead to action under our [DISCIPLINARY POLICY] or, for contractors and suppliers, under their contract. Our response will be proportionate, and clause 13.3 applies to honest mistakes reported promptly.
- 15.2 [IT] may suspend access to an AI tool or agent while an incident is investigated.

16. Legal and regulatory context

- 16.1 This section summarises the main rules relevant to our use of AI as at [DATE OF LAST REVIEW]. It is not legal advice. [DATA PROTECTION LEAD OR LEGAL] keeps it up to date.
- 16.2 **United Kingdom.** The UK GDPR and the Data Protection Act 2018 apply whenever AI tools process personal data. The Data (Use and Access) Act 2025 replaced the UK GDPR's rules on automated decision-making with new Articles 22A to 22D from 5 February 2026 (clause 7.4). The ICO publishes guidance on AI and data protection, which it is reviewing following that Act. (Under the same Act, the Information Commission replaced the Information Commissioner on 30 September 2026; its office is still known as the ICO.) There is no general AI statute in the UK at the time of writing. Regulators such as [THE FCA, THE CMA, OFCOM OR YOUR SECTOR REGULATOR] apply existing law to AI, and equality law (the Equality Act 2010 in Great Britain, separate legislation in Northern Ireland) applies to AI-assisted decisions about people.
- 16.3 **European Union.** The EU AI Act applies to providers placing AI systems on the EU market wherever they are based, to organisations using AI systems that are established in the EU, and to providers and

users outside the EU where the output of an AI system is used in the EU. Its bans on certain practices and its AI literacy duty have applied since 2 February 2025, and most other provisions, including the transparency duties in Article 50, since 2 August 2026. Regulation (EU) 2026/1744, in force since 27 July 2026, moved the obligations for the high-risk uses listed in Annex III, which include recruitment, people management and credit scoring, to 2 December 2027, and those for AI in regulated products to 2 August 2028. The EU GDPR may also apply, for example if we offer goods or services to people in the EU.

16.4 United States. AI rules vary by state and city, are changing quickly, and in some cases have been challenged in court. Examples include New York City’s Local Law 144, which since July 2023 has required bias audits and candidate notices for automated employment decision tools, and Colorado’s SB 26-189, signed in May 2026, which replaces the state’s 2024 AI Act and applies to automated decision-making technology used in consequential decisions, such as employment, lending, housing and insurance, from 1 January 2027. Regulators also apply existing consumer protection, anti-discrimination and sector rules to AI.

16.5 Sector rules. [ADD THE REGULATORS AND RULES THAT APPLY TO YOUR SECTOR, FOR EXAMPLE IN FINANCIAL SERVICES, HEALTHCARE OR LEGAL SERVICES.]

GUIDANCE FOR THE POLICY OWNER • DELETE BEFORE PUBLISHING

We checked every date in this section against the sources in Appendix G on 1 October 2026. Re-check them at each review: several were changed or postponed in 2025 and 2026, and further changes have been proposed. If you operate only in the UK and sell nothing into the EU, keep 16.3 short. If you have US staff or customers, ask counsel which state and city laws apply to you.

17. Review and version control

17.1 The AI Policy Owner reviews this policy at least every [SIX MONTHS], and sooner after a serious incident, a significant change in the law, the adoption of a new category of AI tool (for example the first agents with access to production systems), or a significant change to the business, such as entering a new market.

17.2 The registers in Appendix A are living documents. The AI Register Keeper updates them as decisions are made, and changes to the registers don’t require this policy to be re-approved.

17.3 Version history:

VERSION	DATE	AUTHOR	CHANGES
[1.0]	[DATE]	[NAME]	First version, adapted from the Kiseki Labs AI Acceptable Use Policy template

Approved tools register

The registers are published at [\[LINK\]](#) and kept by the AI Register Keeper. The example rows are illustrations only: replace them with your own decisions.

A1. Approved tools

TOOL AND PLAN	APPROVED USES	HIGHEST DATA CLASS	PERSONAL DATA	OWNER AND REVIEW DATE	CONDITIONS
EXAMPLE: [ASSISTANT NAME] , enterprise plan	Drafting, summarising, analysis, research	Amber	Yes, DPA signed [DATE]	[IT OWNER] , [DATE]	Company SSO, all staff. No training on our data, confirmed in contract. Connectors off except [DOCUMENT STORE] .
EXAMPLE: [CODING ASSISTANT] , business plan	Writing and reviewing code in [APPROVED REPOSITORIES]	Amber	No	[HEAD OF ENGINEERING] , [DATE]	Company SSO, engineering only. Public code filter on. No production credentials.
EXAMPLE: [MEETING NOTETAKER] in [VIDEO PLATFORM]	Summaries of internal meetings	Amber	Yes, DPA signed [DATE]	[IT OWNER] , [DATE]	Opted-in teams only. Announce at the start. External meetings only with agreement.
EXAMPLE: Free and personal tiers of any assistant	Green information only	Green	No	Not applicable	Not approved for company information.

A2. Not approved

TOOL	REASON	USE INSTEAD
[TOOL]	[FOR EXAMPLE: TRAINS ON INPUTS AND NO DATA PROCESSING AGREEMENT IS AVAILABLE]	[APPROVED ALTERNATIVE]

A3. Agents and automations register

One entry per agent or automation (clause 8.2). Copy the blank entry for each new one.

FIELD	EXAMPLE: INVOICE TRIAGE AGENT
PURPOSE	<i>Sorts incoming supplier invoices and drafts queries about missing information</i>
OWNER	[NAME], Finance
CAN READ	Accounts payable inbox; supplier list
CAN DO	Draft replies; tag invoices in [FINANCE SYSTEM]
NEEDS HUMAN APPROVAL FOR	Sending any email; any change to payment details
CREDENTIALS	Its own service account, limited to the accounts payable folder
SPENDING LIMIT	[AMOUNT] a month, hard cap
WHERE LOGS ARE KEPT	[LOGGING TOOL], for [90 DAYS]
HOW TO STOP IT	Disable the service account in [IDENTITY PROVIDER]
REGISTERED ON, REVIEW BY	[DATE], [DATE]
FIELD	[AGENT NAME]
PURPOSE	
OWNER	
CAN READ	
CAN DO	
NEEDS HUMAN APPROVAL FOR	
CREDENTIALS	
SPENDING LIMIT	
WHERE LOGS ARE KEPT	
HOW TO STOP IT	
REGISTERED ON, REVIEW BY	

GUIDANCE FOR THE POLICY OWNER • DELETE BEFORE PUBLISHING

Seed the register from what's already in use. Expense claims, SSO and OAuth grants, browser extensions on managed devices and a short no-blame survey tend to turn up more tools than anyone expects. Then make one decision for each: sanction it, replace it or retire it. Most teams keep these registers in a spreadsheet or their IT service desk tool; the columns and fields here are the ones to use.

Data traffic–light quick reference

Print it, pin it up, or put it in the intranet sidebar.

LIGHT

WHAT IT COVERS

WHERE IT MAY GO

 Green	Public information, and general questions with no company or personal details	Any AI tool, including personal accounts
 Amber	Internal and confidential information: documents, emails, code, financials, contracts, and customer or staff personal data	Approved tools only, on your company account, within the data class the register allows
 Red	Secrets and credentials, card data, special category and criminal records data, privileged material, inside information, security details, and anything a client contract restricts	Only a system approved for that specific information and purpose. Secrets, credentials and card data: never

Before you paste, upload or connect anything, ask

- 1 Which light is this information?
- 2 Is this tool approved for that light? Check the register at [\[LINK\]](#).
- 3 Does the tool need all of it, or can I take out names and details first?

Never, in any AI tool

Passwords, API keys, tokens, private keys and payment card numbers. If one slips through, rotate it and report it to [\[CHANNEL\]](#).

Not sure?

Treat it as the higher class and ask [\[AI POLICY CONTACT\]](#).

The policy on a page

We want you to use AI to do better work. These are the rules that keep it safe. The full [COMPANY NAME] AI Acceptable Use Policy is at [LINK].

Five principles

- **Use approved tools for company work.**
- **Match the data to the tool.**
- **A human owns every output.**
- **Be open about AI use** where it matters to the people affected.
- **Report problems early.**

Do

- Use the tools on the register at [LINK], with your company account.
- Check the traffic-light before you paste, upload or connect anything.
- Check AI output before you use it. If you send it, you own it.
- Get a second pair of eyes on anything where being wrong is expensive.
- Tell people when they're dealing with AI, and announce notetakers in meetings.
- Ask for new tools through [ROUTE]. We aim to decide within [10] working days.
- Report incidents and near misses to [CHANNEL]. Honest mistakes are not punished.

Don't

- Put passwords, keys, tokens or card numbers into any AI tool. Ever.
- Use personal or free accounts for anything except Green information.
- Let AI screen candidates or make decisions about people without approval.
- Create images, audio or video of real people without their consent.
- Install AI extensions, plug-ins or connectors that aren't on the register.
- Run an agent that acts on our systems without registering it.
- Use AI to get around a control, a block or this policy.

The traffic-light

● Green	Public information	Any tool
● Amber	Internal and confidential, including personal data	Approved tools only
● Red	Secrets, card data, special category data, privileged and other restricted information	Only systems approved for it. Secrets: never

Who to contact

Questions	[AI POLICY CONTACT]
Report an incident or near miss	[CHANNEL]
Approved tools register	[LINK]
Policy owner	[NAME, ROLE]

Tool request form

Send the completed form to [EMAIL ADDRESS OR TICKET QUEUE]. The approvers aim to reply within [10] working days. Until a tool is approved, use it with Green information only.

Part 1: Request

QUESTION	ANSWER
Your name, team and the date	
Tool name and provider	
Website	
Plan or tier requested	
Is this a new tool, a new AI feature in an approved tool, or a new connection to another system?	

Part 2: Use and data

QUESTION	ANSWER
What will you use it for, and what problem does it solve?	
Who will use it, and how many people?	
Highest data class it will handle (Green, Amber or Red)	
Will it process personal data? Any special category data?	
Will it connect to other systems, such as email, drives, the CRM or code? Which ones?	
Will it take actions on its own, such as sending, changing, buying or running code?	
Could an approved tool do this instead? If not, why not?	
Cost, pricing model and proposed budget owner	
Trial plan (Green information only until approved)	

Part 3: Approver assessment

QUESTION (CLAUSE 4.3)	YES, NO OR UNKNOWN	NOTES
(a) Training on our data excluded by contract		
(b) Retention acceptable, or zero retention available		
(c) Region and transfer safeguard acceptable		
(d) Business terms in place, with a DPA if personal data is involved		
(e) Single sign-on and central account management		
(f) Usage and audit logs available to administrators		
(g) Sub-processors known, with notice of changes		
(h) Admin controls for connectors, sharing and agent actions		
(i) Security assurance, such as ISO/IEC 27001 or SOC 2 Type II		
(j) Data export and deletion on exit		
(k) Spending limits and alerts available		

Part 4: Decision

ITEM	RECORD
Decision: approved, approved with conditions, or not approved	
Highest data class approved	
Approved uses	
Conditions	
Alternative offered, if not approved	
Review date	
Approved by (names and dates)	
Added to the register on	

Rollout checklist for the policy owner

The first 30 days, week by week. Owners in brackets are suggestions.

Week 1: Find out what's already in use

- ☐ Name the AI Policy Owner, the approvers and the register keeper (section 3).
OWNER [EXECUTIVE SPONSOR]

- ☐ Pull AI subscriptions from expense claims and card statements.
OWNER [FINANCE]

- ☐ List AI apps with single sign-on or OAuth access to [MICROSOFT 365 OR GOOGLE WORKSPACE].
OWNER [IT]

- ☐ List AI browser extensions on managed devices, and AI features switched on in software you already pay for.
OWNER [IT]

- ☐ Inventory the API keys issued for AI providers, and scan code repositories for keys that shouldn't be there.
OWNER [ENGINEERING]

- ☐ Run a short, no-blame survey: what do you use, and what does it save you?
OWNER [AI POLICY OWNER]

Week 2: Decide

- ☐ Triage the top 20 or so tools by usage and data sensitivity, and decide for each: sanction, replace or retire.

- ☐ Map your data classification onto the traffic-light (section 5).

- ☐ Fill in the placeholders, and collect “expensive to be wrong” lists from department heads (clause 7.3).

- ☐ Get the list of clients with AI restrictions from Legal (clause 10.1).

- ☐ Send the draft for legal and data protection review.

Week 3: Publish

- ☐ Approve the policy at [EXECUTIVE COMMITTEE].

☐ Publish the policy, the register and the policy on a page (Appendix C) on [INTRANET].

☐ Open the tool request route and the near-miss channel.

☐ Brief line managers before staff hear about it.

Week 4: Train and enforce

☐ Launch the training and collect acknowledgements (section 14).

☐ Switch on what you can enforce: SSO for approved tools, spending alerts, admin controls, and blocks on unapproved services where an approved alternative exists.

☐ Register any agents already running (section 8).

☐ Book the first policy review for [DATE] and the first quarterly report to [EXECUTIVE COMMITTEE].

Day 30: Check it's working

☐ The share of staff who have acknowledged the policy.

☐ The number of tool requests, and the median time to a decision.

☐ The number of near misses reported. Zero reports in a month is a warning sign.

GUIDANCE FOR THE POLICY OWNER • DELETE BEFORE PUBLISHING

The first two weeks matter most. A policy written without knowing what's already in use tends to ban the tools people rely on, and they'll keep using them anyway.

Glossary

TERM	MEANING
AI agent	An AI system that can take actions, such as sending messages, changing records or running code, as well as producing content.
AI tool	Any software that uses machine learning to generate, summarise, transform, classify or act on content (clause 2.2).
Approved tool	A tool listed as approved in Appendix A, used within its approved data class and uses.
Connector	A link that lets an AI tool read from or act in another system, such as a mailbox, drive, CRM or code repository. Model Context Protocol (MCP) servers are one common type.
Data processing agreement (DPA)	A contract that data protection law requires when a supplier processes personal data on our behalf.
Data protection impact assessment (DPIA)	A structured assessment of the risks of processing personal data, required for processing likely to be high risk.
Deepfake	An AI-generated or altered image, audio or video that shows real people, places or events and could be mistaken for genuine.
Hallucination	A confident but incorrect or invented output from an AI model.
Meaningful human involvement	Review by a person with the authority, information and time to change the outcome.
Model provider	The company whose AI model powers a tool. It may not be the company that sells you the tool.
Open-weight model	An AI model whose weights are published, so it can be run on your own hardware.
Personal data	Information relating to an identified or identifiable living person.
Prompt	The instructions, questions and material you give an AI tool.
Prompt injection	Hidden or malicious instructions in content an AI tool reads, intended to make it ignore its instructions, leak information or take unwanted actions.
Shadow AI	AI tools, features or accounts used for work without the organisation's approval or oversight.
Single sign-on (SSO)	Logging in to a tool with your company account, so access is managed centrally.
Special category data	Personal data about racial or ethnic origin, political opinions, religious or philosophical beliefs, trade union membership, genetic data, biometric data used to identify someone, health, sex life or sexual orientation.
Sub-processor	Another company a supplier uses to process our data, such as the model provider behind an AI tool.
Token	The unit AI providers use to measure, and charge for, the text a model reads and writes.
Training on inputs	When a provider uses the content you give its tool to improve its models.
Zero data retention	An arrangement in which the provider does not store prompts or outputs after processing them.

References

The statements about law and regulation in this template, and the incidents mentioned in the guidance notes, were checked against these sources on 1 October 2026. Re-check them at each review.

Law, regulation and guidance

- 1 European Commission, *AI Act: regulatory framework for AI* (application timeline, and the AI Omnibus in force since 27 July 2026). digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai
- 2 Regulation (EU) 2026/1744 (Digital Omnibus on AI), Official Journal of the EU (published 24 July 2026, in force 27 July 2026). eur-lex.europa.eu/eli/reg/2026/1744/oj
- 3 EU AI Act, Article 2 (scope). artificialintelligenceact.eu/article/2
- 4 EU AI Act, Article 5 (prohibited practices, including emotion recognition in the workplace). artificialintelligenceact.eu/article/5
- 5 EU AI Act, Annex III (high-risk uses, including employment and creditworthiness). artificialintelligenceact.eu/annex/3
- 6 Freshfields, *EU AI Act unpacked #34: the final Digital Omnibus on AI* (new high-risk dates; Article 50 timing). freshfields.com
- 7 Law and Technology, *AI literacy: the Digital Omnibus rewrites Article 4 of the AI Act* (30 July 2026). lawandtechnology.eu
- 8 Data (Use and Access) Act 2025, section 80 (Articles 22A to 22D of the UK GDPR). legislation.gov.uk/ukpga/2025/18/section/80
- 9 DLA Piper, *UK: Commencement of the data protection provisions in the Data (Use and Access) Act* (February 2026). privacymatters.dlapiper.com
- 10 ICO, *The Data (Use and Access) Act 2025: what does it mean for organisations?* ico.org.uk
- 11 ICO, *Guidance on AI and data protection* (marked as under review following the Data (Use and Access) Act). ico.org.uk
- 12 ICO, *Our plans for new and updated guidance: technology* (automated decision-making guidance due winter 2026). ico.org.uk
- 13 ICO, *Personal data breaches: a guide* (72-hour reporting). ico.org.uk
- 14 UK Government, *Information Commission succeeds the ICO as UK's data protection regulator* (30 September 2026). gov.uk
- 15 Lewis Silkin, *AI judgment day on the horizon, while UK lawmakers play catch up* (14 September 2026). lewissilkin.com
- 16 Colorado General Assembly, *SB26-189: Automated Decision-Making Technology* (signed 14 May 2026). leg.colorado.gov/bills/sb26-189
- 17 McDermott, *Colorado AI law in flux: comprehensive replacement bill signed after federal court blocks predecessor's enforcement*. mcdermottlaw.com
- 18 Littler, *New York City adopts final regulations on use of AI in hiring and promotion* (enforcement from 5 July 2023). littler.com

Incidents mentioned in the guidance notes

- 19 Hugging Face, *Security incident disclosure, July 2026* (16 July 2026). huggingface.co/blog/security-incident-july-2026
- 20 Dwarkesh Podcast, *Ajeya Cotra: inside the OpenAI agent swarm that hacked Hugging Face* (1 September 2026). dwarkesh.com/p/ajeya-cotra
- 21 Daniel Trugman (Requesty), talk at the London Software Guild meetup, 24 June 2026, describing a retrospective scan of traffic through Requesty's own AI gateway.

About this template

Kiseki Labs helps mid-market and enterprise organisations in the UK and US get AI into production and keep it there. We start with evidence and come to architecture second.

This template grew out of the four questions a risk team asks about any AI system, usually later than it should: how do you know the output is correct, where does our data go, what happens when it's wrong, and what happens when the people who built it leave. A good acceptable use policy answers the second and third questions for the tools your people already use.

NEXT

A second opinion before you publish.

Thirty minutes on your policy, your register or the agents you're about to deploy, and a straight answer on what needs to change.

[REQUEST AN EXECUTIVE BRIEFING → KISEKILABS.COM/BOOK-A-CALL](https://kisekilabs.com/book-a-call)

DELETE THIS PAGE BEFORE YOU PUBLISH YOUR POLICY.