Anyway Ventures

UK BUSINESSES · HR, DPO AND LEADERSHIP · 2026 EDITION

AI acceptable use *policy template*.

A complete, editable AI acceptable use policy for UK businesses, with drafting notes on every clause, a traffic-light data table, an approved-tools register, request and incident forms, a one-page staff summary and add-ons for regulated sectors. Checked against UK GDPR as amended by the Data (Use and Access) Act 2025, current ICO guidance and the EU AI Act as amended in July 2026.

WORD TEMPLATE

PDF GUIDE

AUTHOR

Brannan Coady, Anyway Ventures

EDITION

Version 1.0 · October 2026

ONLINE, WITH DOWNLOADS

anywayventures.com/resources/ai-acceptable-use-policy-template

CONTACT

brannan@anywayventures.com

Free to use, adapt and share with attribution · Not legal, tax or investment advice

CONTENTS

In this guide

1. Executive summary
2. How to use this pack
3. Adopting the policy in your business
4. The UK legal context
5. Data classification: the traffic-light table
6. The policy text, with drafting notes
7. Appendix A: approved AI tools register
8. Appendix B: AI use-case request and approval form
9. Appendix C: AI incident report form
10. Appendix D: one-page staff summary
11. Manager's FAQ
12. Sector add-ons
13. Worked example
14. Common mistakes
15. Glossary
16. Sources and further reading
17. About this resource

About Anyway Ventures

WHO IT IS FOR

CEOs, COOs, HR directors, DPOs, heads of IT, compliance leads

ALSO AVAILABLE

Word template, PDF guide at anywayventures.com/resources/ai-acceptable-use-policy-template

KEY TAKEAWAYS

The short version

1. No UK law requires a document called an AI policy, but UK GDPR, the Equality Act 2010, client contracts and professional rules already apply to how staff use AI tools.

2. Approve an AI tool before writing rules, because 71% of UK employees had used unapproved consumer AI tools at work and 28% said their employer offered no approved option (Microsoft UK, October 2025).

3. A green, amber and red data table with examples from each team does most of the work in an AI policy, turning a judgement about what is confidential into a quick lookup.

4. Since 5 February 2026 the Data (Use and Access) Act 2025 has replaced UK GDPR Article 22 with Articles 22A to 22D, but a trained person should decide on hiring, pay, performance, discipline and credit.

5. When adopting an AI policy, train staff within 30 days of launch and keep a record, name a policy owner, and set a first review date no more than six months after approval.

This is a complete AI acceptable use policy template for UK businesses, with the guidance needed to adopt it: drafting notes on every clause, a traffic-light data table, an approved-tools register, a use-case request form, an incident report form, a one-page staff summary, a manager's FAQ and add-ons for financial services, legal, healthcare, education and recruitment. It is written for owners, COOs, HR directors, data protection leads and heads of IT in businesses of roughly 20 to 2,000 staff. The editable Word template contains the policy and all the forms, with every blank in [square brackets].

1. Executive summary

1. **No UK law requires a document called an AI policy, but plenty of law already applies.** UK GDPR covers any personal data staff paste into an AI tool, the Equality Act 2010 covers decisions about people, and client contracts and professional rules cover confidential material. A short policy, an approved-tools list and a training record are the simplest way to show you have taken appropriate measures.
2. **Approve a tool before you write rules about tools.** Microsoft's UK research, run by Censuswide with 2,003 employees in October 2025, found 71% had used unapproved consumer AI tools at work and 28% said their employer offered no approved option. A policy that only prohibits moves use onto personal phones.
3. **The data table does most of the work.** "Do not input confidential information" leaves every employee to decide what confidential means. A green, amber and red table with examples from each team turns a judgement into a lookup that takes seconds.
4. **People decide about people.** Since 5 February 2026 the Data (Use and Access) Act 2025 has replaced UK GDPR Article 22 with Articles 22A to 22D, which allow more solely automated decisions, with safeguards. The ICO's 2026 draft guidance and its recruitment report make clear that a rubber-stamp review does not count as human involvement. For most businesses the simple rule is that AI may draft and sort, but a trained person decides on hiring, pay, performance, discipline and credit.
5. **Cover the three risks most templates miss.** AI note-takers recording meetings, AI agents with access to email, files and finance systems, and voice or video impersonation used to push through payments. Clauses 5, 6 and 12 deal with each.
6. **Make it operational, not aspirational.** A policy without a tools register, a request route and an incident route is a statement of values. The Word template includes all three as appendices, ready to use on day one.
7. **Train within 30 days and keep a record.** The ONS reported in July 2026 that around 35% of UK businesses with ten or more employees use AI, but only 11% say more than half their workforce has had AI-related training. The same training record is your evidence under Article 4 of the EU AI Act if you sell into the EU.
8. **Name an owner and a review date no more than six months out.** Provider terms, product features and ICO guidance are all moving. The ICO's main AI guidance page says it is under review because of the Data (Use and Access) Act.

2. How to use this pack

Who should use it. The person who will own the policy (usually the COO, the head of operations or the data protection lead), working with HR, IT or security, and whoever holds client contracts. In a business under 50 staff that may be two people. The CEO or managing director signs it off.

When. Before, or at the same time as, rolling out a company-wide AI tool. If staff already use AI without rules, do it now: the amnesty in clause 11.4 brings existing use into the open. It also belongs in any AI audit or governance review, and in the first 100 days after an acquisition.

How long it takes. Five working days to draft and agree, if the drafting group is small. Four to eight weeks to roll out, train and check adoption.

What you need before you start.

- ☐ A list of AI tools in use: licences, AI features switched on inside existing software, browser extensions and expense claims.
- ☐ The results of a short anonymous staff survey: which AI tools do you use for work, and for what? The shadow AI survey in our AI readiness scorecard is ready to use.
- ☐ The tool or tools you intend to approve, and their business-plan data terms.
- ☐ Your client contracts, NDAs and supplier terms that restrict where information can go.
- ☐ Your existing data protection documents: records of processing, privacy notices, DPIA template.
- ☐ Your disciplinary, information security and bring-your-own-device policies, so the new policy refers to them rather than repeating them.

How the Word template maps to this guide.

WORD TEMPLATE PART	WHERE THE GUIDANCE IS	WHAT TO DO WITH IT
Document control table	Adopting the policy	Fill in owner, approver, dates
Clauses 1 to 15	The policy text, with drafting notes	Replace placeholders, delete what does not apply
Appendix A: approved tools register	Appendix A	Fill in before launch; keep live
Appendix B: use-case request form	Appendix B	Publish with the policy
Appendix C: incident report form	Appendix C	Publish with the policy
Appendix D: one-page staff summary	Appendix D	Print, pin, use in training
Appendix E: sector clauses	Sector add-ons	Insert only those that apply

> **Tip.** Delete before you add. A 25-person firm with no agents and no client meetings can remove clauses 5 and 6 entirely. Every clause you keep is a rule someone has to follow and someone has to enforce.

3. Adopting the policy in your business

Who does what

ROLE	DURING ADOPTION	AFTER LAUNCH
CEO or managing director	Approves the policy and the first tools register; announces it	Signs off each review; receives the incident summary
Policy owner (often COO)	Runs the drafting group; owns the timetable	Keeps the register, answers requests, runs reviews
Data protection lead or DPO	Reviews data rules, DPIA needs, provider terms	Assesses incidents for ICO reporting; approves red-data uses
HR	Checks fit with contracts, disciplinary and monitoring rules; plans training	Keeps training records; handles conduct issues
IT or security	Confirms tools can be configured as the policy says	Enforces the approved list; approves connections and agents
Legal or client contracts owner	Checks client confidentiality and IP terms	Advises on disclosure to clients
Line managers	Read the draft and raise practical objections	Answer staff questions; spot and escalate misuse
Board (PE-backed or larger firms)	Notes the policy and the risk position	Receives a short report at each review

Sign-off route

1. **Draft** from the template, with the policy owner holding the pen.
2. **Written review** by the data protection lead, HR, IT or security, and legal or contracts. One round of comments, in writing, within three working days.
3. **Practical read-through** by four or five staff from different teams. Ask one question: is there anything here you could not follow tomorrow? Their answers are usually more useful than the specialists'.
4. **Approval** by the CEO or managing director, recorded in the document control table.
5. **Board notification** in a PE-backed business, or where the board has asked for it. It answers several of the questions boards should ask about AI.

Consultation

Most AI acceptable use policies are non-contractual workplace rules, like an IT or social media policy, so the employer can introduce and change them without agreeing a variation to each contract. Say so in the policy (clause 15.3 does). Check three things before assuming no formal consultation is needed:

- **Collective agreements and forums.** If you recognise a trade union or have an information and consultation agreement or employee forum, check whether new workplace rules or technology changes must go to it first.

- **Monitoring.** If you will monitor how staff use AI tools beyond ordinary IT logs, the ICO's monitoring workers guidance expects you to be transparent, to be able to justify it, and to carry out a data protection impact assessment where monitoring is likely to result in high risk (the guidance is itself under review following the Data (Use and Access) Act). Consulting staff is good practice and helps show the monitoring is proportionate.
- **AI that affects how staff are managed.** If an AI tool will be used to allocate work, assess performance or monitor productivity, that is a bigger change than an acceptable use policy and needs its own assessment and consultation. If you have EU operations, the EU AI Act separately requires deployers of high-risk systems at work to inform workers' representatives and affected workers first.

Even where nothing is legally required, a two-week window for comments from staff costs little and surfaces the uses you did not know about.

Rollout plan

WHEN	ACTION	OWNER
Week 0	Policy approved; tools register completed; business accounts provisioned	Policy owner, IT
Week 0	CEO announcement: why, what is approved, the 30-day amnesty, training dates	CEO
Week 1	Policy, register, staff summary and forms published on the intranet; staff summary in every induction pack	HR
Weeks 1 to 4	Training for all users, by team, on their own work	HR, trainer
Weeks 1 to 4	Managers' briefing using the manager's FAQ	Policy owner
Week 4	Amnesty closes; disclosed tools reviewed through the request route	Policy owner
Week 4	Personal-account access to unapproved AI sites blocked or warned on company devices, if IT can do so	IT
Week 6	Repeat the anonymous usage survey; compare with the baseline	Policy owner
Week 8	Short report to CEO: training completion, requests, incidents, survey change	Policy owner
Month 6	First formal review	Policy owner, drafting group

Training

Training is where a policy turns into behaviour. People do not read handbooks, and the ones who do cannot turn a general rule into a decision at 4.45pm when a supplier spreadsheet needs summarising. The approach that works is teaching the rules while people use the tools on their own tasks, so the traffic-light table becomes a habit. How to train staff on AI sets out how to structure it.

GROUP	MINIMUM CONTENT	SUGGESTED LENGTH
All users	What the tools do well and badly; the data table with examples from their team; checking output; the incident route	60 to 90 minutes, plus practice on real tasks

GROUP	MINIMUM CONTENT	SUGGESTED LENGTH
Line managers	All of the above, plus the manager's FAQ, how to handle disclosed shadow use, when to escalate	Additional 45 minutes
HR, recruitment and anyone deciding about people	Meaningful human review; bias; Equality Act; the ADM rules	Additional 60 minutes
Finance and anyone approving payments	Impersonation and call-back rules; checking AI figures	Additional 30 minutes
Approvers (policy owner, DPO, IT)	Assessing tools and provider terms; DPIA triggers; incident triage	Half a day
Agent and automation owners	Access scoping, approvals, logging, prompt injection, switching off	Half a day

Record who completed what, when and on which tools. If you have EU exposure, add a column for what each person can now do (for example "built a supplier summary tool with a human check"). That column is what makes the record credible.

Review cycle

Review at a fixed date no more than six months after approval, then every six to twelve months. Review early when any of these happens:

- ☐ A new AI tool is approved, or an AI feature is switched on in existing software.
- ☐ A provider changes its data use, retention or training terms.
- ☐ A reportable incident, or three or more minor incidents of the same kind.
- ☐ The ICO publishes final guidance on AI or automated decision-making.
- ☐ The business starts selling into, or operating in, the EU.
- ☐ The business starts using AI in recruitment, credit, pricing for individuals or another decision about people.
- ☐ An acquisition brings in staff with different tools and habits.

At each review, look at five measures: the percentage of users trained, the number of requests and the median time to decide them, incidents by type, the change in the anonymous usage survey, and any register entries past their review date.

4. The UK legal context

The UK has no AI Act. What applies is the law you already follow, applied to a new kind of tool. This section is accurate to late October 2026 and is not legal advice. Each part ends with the clauses that respond to it.

UK GDPR and the Data Protection Act 2018

If someone pastes a customer complaint, a CV or a payroll export into an AI tool, the business is processing personal data. The ordinary rules apply: a lawful basis (Article 6, and Article 9 for special category data), transparency in your privacy notices (Articles 13 and 14), a written processor contract with the provider (Article 28),

appropriate security (Article 32), lawful international transfers (Chapter V), and a data protection impact assessment where processing is likely to result in high risk (Article 35). Personal data breaches likely to result in a risk to people must be reported to the ICO without undue delay and within 72 hours of becoming aware (Article 33), as the ICO's breach guidance explains.

The practical points for an acceptable use policy are that business plans with processor terms are materially different from consumer accounts, that you need to know where data is stored and whether it is used to train models, and that someone must decide whether a new use needs a DPIA. Pseudonymised data is still personal data if you can re-identify it; only properly anonymised data falls outside UK GDPR.

Clauses: 3, 4, 5, 11, 14.

The Data (Use and Access) Act 2025

The Act amends UK GDPR rather than replacing it. Three changes matter for an AI policy.

Automated decision-making. Section 80 replaced Article 22 with Articles 22A to 22D, in force from 5 February 2026 under SI 2026/82. A decision is "based solely on automated processing" if there is no meaningful human involvement in taking it, and it is significant if it has a legal or similarly significant effect. Solely automated significant decisions are now allowed more widely, but where one is made the controller must have safeguards that let the person be told about it, make representations, obtain human intervention and contest it (Article 22C). Solely automated significant decisions based on special category data are restricted to explicit consent or limited legal grounds (Article 22B), and no solely automated significant decision can rely on the new "recognised legitimate interests" lawful basis.

What "meaningful" means. The ICO consulted on draft guidance on automated decision-making and profiling from 31 March to 29 May 2026. As summarised by law firms including Brodies and Covington, the draft expects human involvement to be active, carried out by someone with the authority and competence to change the outcome, and to happen before the decision takes effect. Its Recruitment rewired report, drawing on evidence from over 30 employers and published the same day, found that many employers who believed a human was in the loop did not have meaningful human involvement. Final guidance is expected later in 2026, so check for it at your next review.

Data protection complaints. From 19 June 2026, section 103 requires every controller to facilitate complaints (for example with an electronic form), acknowledge them within 30 days and respond without undue delay. Complaints about AI-assisted decisions will come through this route, so the incident process in clause 14 should connect to it.

Clauses: 7, 8, 14. Sector add-on: recruitment.

The Equality Act 2010

The Equality Act 2010 applies whatever tool produced a decision. An AI screening tool that disadvantages candidates sharing a protected characteristic can amount to indirect discrimination (section 19), a failure to make reasonable adjustments for disabled candidates or staff (section 20) can arise from an inflexible automated process, and employers are liable for what employees and agents do in the course of employment (section 109). The ICO's 2024 audits of AI recruitment tool providers found tools that inferred gender and ethnicity from names, and others that allowed filtering by protected characteristics. Public bodies, including state schools, also have the public sector equality duty in section 149.

Clauses: 7.1(a) and (b), 8.3(c). Sector add-ons: recruitment, education.

ICO guidance on AI

The ICO is the regulator most UK businesses will meet on AI. Its core documents:

- Guidance on AI and data protection: accountability, lawfulness, fairness, security and data minimisation for AI. The page states that because of changes made by the Data (Use and Access) Act, the guidance is under review and may be subject to change.
- The draft ADM and profiling guidance and the Recruitment rewired report, above.
- Monitoring workers: relevant if you monitor AI use, and to any AI tool that monitors staff.

Nothing in the ICO's guidance asks for a document called an AI policy. It asks you to be able to show your measures. The policy, register, DPIAs, training record and incident log together are that evidence.

Confidentiality, contracts and privilege

Client contracts, NDAs, supplier terms and professional duties do not make an exception for AI tools. Common restrictions to look for: information must stay within named systems or jurisdictions, may not be disclosed to sub-processors without consent, or must be deleted at the end of an engagement. Putting it into a tool whose provider retains it, or uses it for training, can breach those terms. In legal work there is the added question of whether disclosure to a tool waives privilege, which the SRA has flagged (see the legal add-on).

Clauses: 4.1 (red data), 9.2, 10.2.

Copyright and intellectual property

Three points are settled enough to write into a policy. First, work created by employees in the course of employment belongs to the employer, with or without AI. Second, the UK still has a special rule for "computer-generated" works with no human author (section 9(3) of the Copyright, Designs and Patents Act 1988); the government's Report on Copyright and Artificial Intelligence, published in March 2026 under section 136 of the Data (Use and Access) Act after a consultation that drew 11,520 responses, signalled that this provision could be removed, but has not legislated, and stepped back from a broad training exception. Third, output can still infringe if it reproduces a substantial part of someone else's work. In Getty Images v Stability AI ([2025] EWHC 2863 (Ch)), the High Court in November 2025 held that the trained model was not itself an infringing copy and rejected most of Getty's claims, but that judgment concerns the developer, not a business publishing what a tool produced.

The policy response is practical: check output that will be published or shipped, especially code, images and long passages; do not feed in licensed material whose terms forbid it; and keep a record of significant human contribution where ownership matters. Some enterprise plans include an IP indemnity for output, usually with conditions such as keeping content filters switched on; record whether yours does in the tools register.

Clauses: 10.

Consumer protection and advertising

Using AI to write fake reviews is now expressly banned. The Digital Markets, Competition and Consumers Act 2024 lists submitting or commissioning fake consumer reviews, and publishing them in a misleading way, among the commercial practices that are always unfair, from 6 April 2025. Ordinary misleading-advertising rules apply to AI-generated marketing in the same way as to any other.

Clauses: 7.1(b), 8.3(a).

The EU AI Act, if you sell into the EU

The EU AI Act (Regulation (EU) 2024/1689) follows the market, not the passport. It applies to providers placing AI systems on the EU market and to deployers whose use, or whose systems' output, is used in the EU. A UK business is most likely to be caught if it sells an AI-enabled product to EU customers, has an EU subsidiary or office, or uses AI to produce work delivered to and relied on in the EU. A purely domestic business is not. When the EU AI Act's AI literacy duty applies to UK businesses covers the scope tests in more detail.

The Act was amended by the Digital Omnibus on AI, Regulation (EU) 2026/1744, which entered into force on 27 July 2026. As amended, the parts that matter for an acceptable use policy are:

PROVISION	WHAT IT REQUIRES	WHEN
Article 4, AI literacy	Providers and deployers must take measures to support the development of AI literacy of staff and others using AI systems on their behalf. The Commission's AI literacy Q&A says it does not require any guaranteed level, and no certificate is needed; an internal record is enough	Applied since 2 February 2025; reworded July 2026; national authorities supervising since 2 August 2026
Article 5, prohibited practices	Bans on, for example, manipulative techniques, social scoring and emotion recognition in the workplace (except for medical or safety reasons)	Since 2 February 2025; two added bans (non-consensual intimate imagery and child sexual abuse material) from 2 December 2026
Article 50, transparency	People must be told when they interact with an AI system; deployers must disclose deepfakes and, in some cases, AI-generated text published to inform the public	From 2 August 2026, with a grace period to 2 December 2026 for marking output from generative systems already on the market
Article 26, deployers of high-risk systems	Human oversight by people with competence, training and authority; keeping logs; informing workers before using high-risk AI at work; informing people subject to decisions	High-risk obligations for Annex III uses (including employment, education and credit) postponed by the Omnibus to 2 December 2027; for AI in Annex I regulated products, 2 August 2028

Commentators including ActuIA have summarised the Omnibus dates; check the consolidated text on EUR-Lex before relying on any single date, because further guidance and national implementing laws are still arriving.

Clauses: 9.3, 13. Sector add-ons: recruitment, education.

Summary: which law drives which clause

LAW OR GUIDANCE	APPLIES WHEN	MAIN CLAUSES
UK GDPR, DPA 2018	Any personal data goes into or comes out of an AI tool	3, 4, 5, 11, 14
Data (Use and Access) Act 2025	AI is used in a significant decision about a person; any data protection complaint	7, 8, 14
Equality Act 2010	AI touches recruitment, employment, services to the public	7, 8
Confidentiality, contracts, privilege	Client or third-party information is involved	4, 9, 10
CDPA 1988 and case law	AI output is published, shipped or sold	10

LAW OR GUIDANCE	APPLIES WHEN	MAIN CLAUSES
DMCC Act 2024	AI is used in marketing, reviews or customer communications	7, 8
EU AI Act as amended	EU market, EU staff or EU-affected people	9, 13, sector add-ons
Sector regulators	FCA, SRA, health, education bodies	Sector add-ons

5. Data classification: the traffic-light table

This table is the heart of the policy. Clause 4 refers to it, and training should spend most of its time on it. Replace the examples with ones from your own teams: people learn the rule from the example, not the definition.

COLOUR	WHAT IT INCLUDES	WHERE IT CAN GO
Green	Public information; published marketing; general knowledge questions; your own writing with no names or identifiers; properly anonymised examples	Any approved tool, including personal and free accounts
Amber	Internal documents and reports; non-sensitive business data; supplier and pricing information; customer names and business contact details; meeting notes not covered by red	Approved business accounts only (Appendix A), within the tool's approved level
Red	Special category data (health, ethnicity, religion, sexual orientation, trade union membership, biometrics and similar); criminal records; HR, disciplinary, grievance and appraisal records; payroll and bank details; passwords, keys and credentials; unreleased financial results and inside information; information an NDA or contract restricts; children's personal data; anything a client has asked to keep within named systems	No AI tool unless the policy owner and data protection lead have approved that specific use in writing, in a specific tool

Examples by team

TEAM	GREEN	AMBER	RED
Finance	Explaining an accounting standard; a formula with dummy numbers	Management accounts; aged debt report; supplier price lists	Payroll; bank details; unreleased results; anything inside information
HR	Drafting a generic job advert; summarising published employment guidance	Anonymised engagement survey themes; organisation charts	CVs and interview notes; sickness records; disciplinary or grievance files; appraisal ratings
Sales and marketing	Rewriting a published case study; brainstorming campaign ideas	CRM notes with business contacts; pipeline reports; draft proposals	Contracts with confidentiality terms; data bought under restrictive licences
Operations	Writing a generic process description	Internal procedures; supplier performance data; incident logs without personal data	Security configurations; access lists; customer data covered by contract restrictions
Customer service	Template replies with no customer details	Individual complaints with business contact details	Vulnerable customer information; health details; payment card data
Technology	Public documentation questions; open-source code	Internal code without secrets, in an approved coding tool	API keys, passwords, tokens; production customer data; security vulnerabilities not yet fixed

Decision rules

1. **If you are not sure of the colour, treat it as red and ask.**
2. **The highest colour wins.** A spreadsheet with one red column is red.
3. **Combinations can change the colour.** Two amber files joined together can identify people or reveal something neither does alone.
4. **Making data greener is allowed.** Remove names, replace them with roles ("Customer A", "the supplier"), round figures and strip identifiers. Pseudonymised data is still personal data if the key exists; only properly anonymised data is green.
5. **The tool's level and the data's colour must both allow it.** Amber data in a tool approved only for green is a breach, even if the tool is a business account.
6. **Output inherits the colour of the input.** A summary of an amber document is amber.

> **Drafting note.** Some businesses add a fourth level, "black" or "prohibited", for data that must never go into any AI tool even with approval, such as payment card numbers subject to PCI DSS or material under a court order. That is a sensible addition in financial services and healthcare. Most businesses can manage with three, provided red requires named, written approval for each use.

6. The policy text, with drafting notes

The policy below is the same text as the Word template. Replace everything in [square brackets]. After each clause, a drafting note explains why it is there and how to make it lighter or stricter. The notes are not in the Word document.

1. Purpose

1.1 This policy explains how [Company name] staff may use artificial intelligence (AI) tools, including generative AI assistants, at work.

1.2 We want people to use AI to do better work faster. This policy sets out how to do that without putting customers, colleagues or the business at risk.

> **Drafting note.** Why it is there: the tone of the first clause tells staff whether this is a permission or a prohibition. Policies that open with risk get read as a ban, and use goes underground. Lighter: keep it as it is. Stricter: if your sector requires it, add "and in line with our obligations to [regulator]" to 1.2.

2. Scope and definitions

2.1 This policy applies to all employees, contractors, temporary staff and anyone else working on [Company name]'s behalf.

2.2 It covers any AI tool used for work, on any device, including tools built into software we already use (for example email, office, video call or CRM software), browser extensions, and tools staff build themselves.

2.3 In this policy, an "AI tool" is any software that generates text, images, audio, video, code or analysis in response to instructions. An "AI agent" is an AI tool that can take actions on our systems, such as sending emails, updating records or making payments, as well as producing content.

2.4 This policy sits alongside our [information security policy], [data protection policy], [bring your own device policy] and [disciplinary procedure]. Where they overlap, follow the stricter rule.

> **Drafting note.** Why it is there: most AI use now arrives inside software you already pay for, switched on by a vendor update, and on personal phones. A scope limited to "company-provided AI tools" misses both. Lighter: a business with no contractors can shorten 2.1. Stricter: regulated firms should extend 2.1 to "and anyone acting as our appointed representative or agent", and add a duty for contractors to follow equivalent rules in their contracts.

3. Approved tools

3.1 The approved AI tools are listed in the approved tools register at Appendix A, which sets out for each tool the plan, who may use it, and the highest data colour it is approved for. Only use these tools for work, and only through accounts provided by [Company name].

TOOL	PLAN	WHO MAY USE IT	APPROVED FOR DATA UP TO
[e.g. Microsoft 365 Copilot]	[Business account]	[All staff]	[Amber]
[e.g. Claude Team]	[Business account]	[Finance and legal]	[Amber]
[Public web tools, free and personal accounts]	n/a	Anyone	[Green only]

3.2 Do not use personal or free accounts of any AI tool for work involving anything other than green data.

3.3 Browser extensions and plug-ins that use AI count as separate tools and need their own approval, even if they come from an approved provider.

3.4 The approved list is maintained by [policy owner] and published at [location].

3.5 Do not change the settings of an approved tool that [IT] has set, including settings that control whether our data is used to train models, how long it is kept, and which other systems it can connect to.

> **Drafting note.** Why it is there: the approved list is what makes everything else workable. "Use AI responsibly" is not a rule; "use Copilot through your work login" is. Lighter: a small firm can keep the three-row table here and skip Appendix A. Stricter: remove the green-only row for personal accounts and block public AI sites on company devices; this is common in financial services and healthcare, but only works if the approved tool is good enough that people do not reach for their phones. If you have not chosen a tool, the comparison of Copilot, ChatGPT and Claude for business sets out how to decide.

4. Data rules

4.1 Before putting any information into an AI tool, check its colour in the table below and the tool's approved level in Appendix A.

COLOUR	WHAT IT INCLUDES	WHERE IT CAN GO
Green	Public information, published marketing content, general questions, anonymised examples you have written yourself	Any approved tool, including public tools
Amber	Internal documents, non-sensitive business data, supplier and pricing information, internal reports, customer names and business contact details	Approved business accounts only, as listed in Appendix A
Red	Special category personal data (health, ethnicity, religion and similar), HR and disciplinary records, payroll and bank details, passwords and credentials, unreleased financial results, information a contract or NDA restricts, anything the client has asked us to keep within named systems	Not into any AI tool unless [policy owner] and [data protection lead] have approved that specific use in writing

4.2 If you are not sure of the colour, treat it as red and ask [policy owner].

4.3 Remove names and identifying details where the task does not need them.

4.4 The highest colour in a document sets its colour, and AI output has the same colour as the information that went in.

> **Drafting note.** Why it is there: this is the clause that prevents the common incident, which is not exotic. It is someone pasting a payroll export or a disciplinary note into a free chatbot because it was quicker than asking. Lighter: keep the table; trim the examples to your own. Stricter: add a black level (see the data classification section), require DPIAs for any approved red use, and add "client-identifiable information is red unless the client has agreed in writing to the use of the specific tool", which is the default position in many professional services firms.

5. AI meeting assistants and recording

5.1 Only use the AI note-taking or transcription tools listed in Appendix A.

5.2 Tell everyone in the meeting before an AI tool records or transcribes it, and stop if anyone objects.

5.3 Do not use AI recording or transcription in HR, disciplinary, grievance, medical or legal meetings, or in client meetings where the client has not agreed, unless [policy owner] has approved it in writing.

5.4 Treat transcripts and AI meeting summaries as at least amber data. Check summaries for accuracy before sharing them, and delete recordings and transcripts after [90 days] unless there is a business reason to keep them.

5.5 Do not let an AI note-taker join meetings you are not attending, and switch off automatic joining for external meetings unless [policy owner] has approved it.

> **Drafting note.** Why it is there: note-takers are now built into most video platforms and are often switched on by default. A transcript of an appraisal or a client negotiation is personal and often confidential data, and an AI summary can attribute words to the wrong person. Lighter: a firm with no external meetings can merge 5.2 to 5.5 into one sentence. Stricter: financial services firms that must record certain communications should make the note-taker's transcript subject to their recording and retention rules (see the financial services add-on); healthcare and legal settings should treat all transcripts as red by default.

6. AI agents and connections to company systems

6.1 Do not connect an AI tool to company systems (email, calendars, file storage, CRM, finance or HR systems) unless IT has approved that connection.

6.2 Every AI agent must have a named owner, a short written description of what it does and what it can access, and its own login with the least access it needs. Do not let an agent use your personal credentials.

6.3 A person must approve an agent's action before it sends anything outside [Company name], changes financial or HR records, or makes or approves a payment.

6.4 The owner keeps a log of what the agent has done, reviews exceptions at least [weekly], and switches the agent off if it behaves unexpectedly.

6.5 Do not give an agent access to confidential information, content from outside the business (such as incoming emails or web pages) and the ability to send messages externally all at once, unless IT has assessed and approved that combination.

> **Drafting note.** Why it is there: agents act, not just answer. An agent that can read incoming email can be instructed by whoever writes the email (prompt injection), and one with your credentials can do anything you can. Clause 6.5 targets the combination security researchers consider most dangerous: private data, untrusted input and a way to send data out. Lighter: if no one is building or connecting agents, replace clause 6 with "Do not connect AI tools to company systems without IT approval." Stricter: require a short risk assessment for every agent, an agent register (Appendix A, part 2), a kill switch tested before go-live, and dual approval for any payment-related action.

7. Prohibited uses

7.1 Do not use AI tools to:

(a) make final decisions about recruitment, pay, performance, discipline or dismissal;

(b) create content that is misleading, discriminatory, harassing or unlawful, including fake reviews or testimonials;

(c) impersonate a real person, or create images, audio or video of a real person without their consent;

(d) get round security controls, licence terms or this policy;

(e) process red data except as allowed in clause 4.1;

(f) give regulated advice (for example legal, financial or medical) to customers without the review required in clause 8;

(g) monitor, profile or assess colleagues, or infer their emotions, health or personal characteristics, unless that use has been approved under clause 11.

> **Drafting note.** Why it is there: a short list of bright lines is easier to remember and enforce than a long list of cautions. 7.1(a) is the most important: it keeps most businesses well clear of the automated decision rules in Articles 22A to 22D. 7.1(g) reflects the ICO's monitoring workers guidance and, for EU operations, the AI Act's ban on emotion recognition at work. Lighter: none of these should be deleted, but 7.1(f) can go if you give no regulated advice. Stricter: add "use AI to generate content for children" in education, "use AI to triage clinical urgency" in healthcare, or "use AI to produce investment recommendations" in financial services, unless approved through the sector process.

8. Human review

8.1 You are responsible for any work you produce with AI, as if you had written it yourself.

8.2 AI tools make mistakes, including confident mistakes about facts, figures and sources. Check anything that will be relied on.

8.3 A named person must review AI output before it is:

(a) sent to a customer, supplier, regulator or the public;

(b) used in a financial figure, report or forecast;

(c) used to support a decision about a person.

8.4 Automated tools built with AI must have a named owner, a short written description of what they do, and a person who reviews exceptions.

8.5 A review only counts if the reviewer has the knowledge and the authority to change the result, has looked at the evidence rather than only the AI's summary, and does it before the output is used.

8.6 Code written with AI tools must go through the same review and testing as any other code before it is deployed.

> **Drafting note.** Why it is there: 8.5 is the clause most templates lack, and it matters more since 2026. The ICO's draft ADM guidance and its recruitment report both make the point that a human who clicks "approve" on every recommendation is not meaningful involvement. Lighter: a firm with no decisions about people can drop 8.3(c) and 8.5. Stricter: require a second reviewer for output going to regulators or courts, keep a review log for 8.3(c), and add sampling: the owner checks [5%] of AI-assisted customer responses each month.

9. Transparency and disclosure

9.1 Tell your manager if AI was used to produce a significant piece of work.

9.2 Tell a customer when AI has materially produced work they are paying for, if their contract requires it, or if they would reasonably expect to know.

9.3 Where customers interact with an AI tool directly (for example a chatbot), make it clear they are not talking to a person and give them a way to reach one.

9.4 Label images, audio or video generated or substantially altered by AI that could be mistaken for real, when we publish them.

> **Drafting note.** Why it is there: disclosure protects the relationship as much as the law. 9.3 and 9.4 match the direction of Article 50 of the EU AI Act, which applies from 2 August 2026 to systems used in the EU, and good practice in the UK. Lighter: a business with no customer-facing AI can drop 9.3. Stricter: professional services firms often add a standard sentence for engagement letters, and some clients now require notice before any AI tool touches their matter; check client terms and put the requirement into 9.2.

10. Intellectual property and confidentiality

10.1 Work you create with AI tools in the course of your job belongs to [Company name], as with any other work.

10.2 Do not put third-party confidential material, or material we licence under restrictive terms, into an AI tool unless its terms allow it.

10.3 Check AI output for content that may copy someone else's work, especially code, images and long passages of text, before you publish or ship it.

10.4 Where ownership of a piece of work matters (for example designs, software or content we sell), keep a short record of the human contribution: the brief, the main choices and the edits.

> **Drafting note.** Why it is there: ownership of AI-assisted work is clear for employees, but whether a purely AI-generated work has copyright protection in the UK is under review, and output can still infringe. 10.4 helps prove authorship if it is challenged. Lighter: a business that publishes little can drop 10.4. Stricter: software businesses should add rules for AI coding tools: licence scanning on AI-written code, no pasting of proprietary code into tools without a no-training commitment, and a record of which repositories allow AI assistance. This matters at exit, when buyers ask about code provenance in technical due diligence.

11. New tools and procurement

11.1 To request a new AI tool, or a new use of an approved tool involving red data, complete the AI use-case request form at Appendix B and send it to [policy owner].

11.2 Before approval, [IT/security] and [data protection lead] will check: the provider's data protection terms, where data is stored, whether our data is used to train models, security controls, and cost. We aim to decide within [10 working days].

11.3 Do not buy AI tools on a company card or expense them without approval.

11.4 If you are already using an AI tool that is not on the approved list, tell [policy owner] within [30] days of this policy being published. You will not be disciplined for use disclosed in that period, unless red data was involved.

11.5 [Policy owner] keeps a log of requests and decisions, with reasons, and publishes the approved list after each change.

> **Drafting note.** Why it is there: a request route with a deadline is the alternative to shadow use. If approval takes three months, people will not wait. The amnesty in 11.4 is the fastest way to find out what is really in use. Lighter: a small firm can replace the form with an email to the owner covering the same questions. Stricter: add security certification requirements (for example ISO/IEC 27001 or SOC 2), an exit plan for data at the end of the contract, and a requirement that the request includes a DPIA screening.

12. Security and impersonation

12.1 AI makes it cheap to fake a voice, a video or a convincing email. Do not treat a familiar voice or face as proof of identity.

12.2 Verify any request to change bank details, make an urgent payment or share credentials by calling the person back on a number we already hold, however senior the person appears to be.

12.3 Report suspected AI-generated phishing or impersonation to [IT/security] straight away.

12.4 Never enter passwords, keys or other credentials into an AI tool, and do not install AI tools or extensions on company devices without IT approval.

> **Drafting note.** Why it is there: in 2024 the engineering firm Arup confirmed it had been the target of a deepfake scam in which an employee in Hong Kong, after a video call with what appeared to be the chief financial officer and colleagues, sent about US$25 million, according to Hong Kong police. The control that stops this is procedural, not technical: a call-back on a number already held. Lighter: none; keep it. Stricter: add dual authorisation above [£10,000], a code word for urgent payment requests, and a rule that no payment instruction is accepted from a video or voice call alone.

13. Training

13.1 Everyone who uses an approved AI tool must complete [Company name]'s AI training within [30] days of being given access, and refresh it every [12] months.

13.2 Training covers what the tools are good and bad at, the data rules in this policy, and how to build and check AI-assisted work in your own role.

13.3 [HR] keeps a record of who has completed training, what it covered and which tools it related to.

> **Drafting note.** Why it is there: a policy nobody has been trained on is just a PDF. The record in 13.3 is also your evidence for Article 4 of the EU AI Act if you are in scope; as amended in July 2026, it asks for measures to support AI literacy, recorded internally, not certificates. Lighter: a small firm can run one session and record attendance. Stricter: role-based modules (see the training table in the adoption section), a short practical assessment, and no access to red-data uses until the relevant module is complete.

14. Incidents

14.1 If you think red data has gone into an AI tool by mistake, an agent has done something it should not have, or AI output has caused an error that reached a customer or a financial figure, report it to [policy owner] and [data protection lead] straight away, within [24 hours], using the form at Appendix C.

14.2 Reporting a mistake quickly will not, by itself, lead to disciplinary action. Hiding one may.

14.3 [Data protection lead] will decide whether the incident is a personal data breach that must be reported to the ICO, which must be done within 72 hours of us becoming aware of it.

14.4 [Policy owner] keeps a log of AI incidents and near misses, and reviews it at each policy review.

> **Drafting note.** Why it is there: the 72-hour clock under UK GDPR Article 33 starts when the business becomes aware, which can be the moment any employee knows. A 24-hour internal deadline leaves the data protection lead time to assess. 14.2 matters more than it looks: people report mistakes when reporting is safe. Lighter: keep all of it. Stricter: add contact details for out-of-hours reporting, a requirement to preserve the conversation and output, and notification to clients where contracts require it within a set time.

15. Ownership and review

15.1 This policy is owned by [name and role].

15.2 It was approved by [CEO name] on [date] and will be reviewed by [date, no more than six months later], or sooner if the law, our tools or ICO guidance change.

15.3 This policy does not form part of any employee's contract of employment, and we may change it at any time. Breaches of this policy may be dealt with under [Company name]'s disciplinary procedure.

15.4 We may monitor the use of approved AI tools through their normal administrative logs to check this policy is being followed, as explained in our [staff privacy notice]. We will not use those logs to assess individual performance without telling you first.

> **Drafting note.** Why it is there: tools change monthly, and a policy without a named owner and date will be out of date by the spring. 15.3 keeps the policy non-contractual so it can be updated quickly. 15.4 makes monitoring transparent, which the ICO's monitoring workers guidance expects. Lighter: drop 15.4 if you will not look at logs. Stricter: add a short annual report to the board and a list of the review triggers from the adoption section.

7. Appendix A: approved AI tools register

Keep one row per tool and plan. The register is the live part of the policy: update it whenever a tool is approved, changed or withdrawn, and publish the current version at the location in clause 3.4.

Part 1: tools

FIELD	WHAT TO RECORD	ILLUSTRATIVE ENTRY
Tool and plan	Product name and the exact plan or licence tier	[Example: Microsoft 365 Copilot, enterprise licence]
Provider and contract	Contracting entity; whether a processor agreement is in place	[Provider name; DPA signed date]
Approved users	Teams or roles	[All staff with an M365 licence]
Highest data colour	Green, amber or red (red only with a named approved use)	[Amber]
Approved uses	What it is approved for	[Drafting, summarising, analysis of internal documents]
Not approved for	Specific exclusions	[HR records; client data under NDA]
Training on our data	Whether the provider uses our inputs to train models, and the setting that controls it	[No, per business terms; setting locked by IT]
Data location and retention	Where data is processed and how long it is kept	[UK or EU; retention per tenant settings]
Connections allowed	Other systems it may access	[SharePoint, Outlook; no finance system]
Note-taker or recording features	Whether enabled, and rules	[Enabled for internal meetings only]
IP indemnity	Whether output is covered, and conditions	[Yes, subject to provider conditions]
DPIA	Reference and date, or "screened, not required"	[DPIA-2026-04, 12 October 2026]
Owner	Named person accountable for the tool	[Name, role]
Approved by and date	Who approved, when	[Policy owner and DPO, date]
Next review	Date	[Six months from approval]

Part 2: AI agents and automations

FIELD	WHAT TO RECORD
Agent name and purpose	[What it does, in one sentence]
Owner	[Named person]

FIELD	WHAT TO RECORD
Systems and data it can access	[List, with access level: read or write]
Actions it can take	[For example: drafts replies; updates CRM fields]
Actions needing human approval	[External sends, financial or HR changes, payments]
Credentials	[Dedicated service account; scope]
Untrusted inputs	[Incoming email, web pages, uploaded files]
Log location and review frequency	[Where; how often reviewed]
How to switch it off	[Steps and who can do it]
Approved by, date, next review	[Names and dates]

8. Appendix B: AI use-case request and approval form

Use this to request a new AI tool, a new use of an approved tool involving red data, a new connection to company systems, or an AI agent. One form per request.

Part 1: completed by the requester

QUESTION	RESPONSE
Your name, role and team	[ ]
What is the task, and how is it done today?	[Describe the work, how often, how long it takes]
What would the AI tool do?	[Draft, summarise, analyse, decide, act]
Which tool and plan? Is it already on the approved list?	[Tool, plan, link to terms]
What data would go in? Highest colour?	[Describe the data; green, amber or red]
Does it involve personal data? Whose?	[Customers, staff, candidates, children, none]
Will output be used in a decision about a person?	[Yes or no; describe]
Will output go outside the business?	[To whom]
Will it connect to any company system or take actions?	[Systems; read or write]
Who will check the output, and how?	[Named reviewer; what they check]
Expected benefit	[Hours saved per month, quality, speed]
Cost	[Licences, set-up, running costs]
Is there a client contract or regulator rule that applies?	[Reference]

Part 2: completed by the reviewers

- ☐ Provider terms reviewed: processor agreement, data location, retention, no training on our data or a setting to stop it.
- ☐ Security reviewed: single sign-on, admin controls, audit logs, certifications.
- ☐ DPIA screening done; full DPIA completed if needed.
- ☐ Decisions about people: meaningful human review designed; Equality Act risks considered.
- ☐ Client contracts and confidentiality checked.
- ☐ Connections and agent permissions limited to what is needed.
- ☐ Training needs identified.
- ☐ Register entry drafted (Appendix A).

DECISION	RESPONSE
Outcome	[Approved / approved with conditions / declined]
Conditions	[For example: amber only; finance team only; review after three months]
Reasons	[Short reasons, recorded in the request log]
Reviewed by and date	[IT/security, data protection lead]
Approved by and date	[Policy owner]

9. Appendix C: AI incident report form

Report within [24 hours] of noticing. If in doubt, report: a near miss is useful information and will not count against you.

QUESTION	RESPONSE
Your name, role, contact details	[ ]
Date and time of the incident, and when you noticed	[ ]
What type of incident?	[Red data entered; wrong tool used; agent acted wrongly; AI error reached a customer or figure; impersonation or phishing; recording without notice; other]
Which tool, account and plan?	[Tool; personal or business account]
What happened?	[Plain description, in order]
What data was involved? Whose? How many people?	[Describe; estimate numbers]
Has the output been sent, published or relied on?	[Where, to whom]
What have you done so far?	[Deleted the conversation, recalled the email, told the customer]

QUESTION	RESPONSE
Have you kept the evidence?	[Screenshots, conversation link, output]

For the data protection lead and policy owner

ASSESSMENT	RESPONSE
Is it a personal data breach?	[Yes or no, with reasons]
Likely risk to individuals	[None, low, risk, high risk]
ICO notification needed? Deadline	[Yes or no; 72 hours from awareness: date and time]
Individuals to be told?	[Yes or no; when]
Clients or other parties to be told under contract?	[Who; deadline]
Provider asked to delete data?	[Request date; confirmation]
Root cause	[Unclear rule, training gap, tool setting, deliberate breach]
Actions to prevent repeat	[Policy change, training, setting, register update]
Closed by and date	[ ]

10. Appendix D: one-page staff summary

Give everyone this page alongside the full policy. It is what people should be able to recite after training.

1. **Use the approved tools, through your work account.** The list is at [location]. Personal accounts are for green data only.
2. **Check the colour before you paste.** Green goes anywhere approved. Amber goes into business accounts only. Red doesn't go in without written approval. Not sure? It's red.
3. **You own the output.** Check facts, figures and sources before anything leaves the building or goes into a number.
4. **People decide about people.** AI can help draft, never decide, on hiring, pay, performance or discipline.
5. **Say so when it matters.** Tell your manager about significant AI-assisted work, and the customer when they'd expect to know.
6. **Ask before you buy or connect.** New tools, note-takers and connections to company systems go through [policy owner].
7. **Check before you pay.** A familiar voice on a call is not proof. Call back on a number you already hold.
8. **Report mistakes fast.** Within [24 hours], to [policy owner]. Owning up quickly is treated well.

Questions: [policy owner name, email]. Request form, incident form and approved list: [location].

11. Manager's FAQ

Managers get the questions first. These are the ones that come up in the first month, with answers consistent with the policy. Adapt the bracketed parts.

"Can I use ChatGPT on my phone for work?" Only for green data: public information, or your own writing with no names or internal details. For anything internal, use [approved tool] through your work account, which has terms that protect our data.

"The approved tool is worse than the one I use at home. Can I keep using mine?" Not for amber or red data. Submit a request (Appendix B) explaining what the other tool does better. If enough people ask, it will be reviewed. Meanwhile, tell [policy owner] about any existing use during the amnesty.

"I've been using an unapproved tool for months. Am I in trouble?" Not if you tell [policy owner] within 30 days of the policy being published, unless red data was involved. Even then, telling us is far better than us finding out. As a manager, thank people who disclose.

"Is a customer's name and email amber or red?" Business contact details are amber. A customer's health, finances as an individual, or anything they told us in confidence is red. If the customer is vulnerable, treat it as red.

"Can I use AI to write performance reviews?" You can use it to tidy your own wording, with names removed. You cannot paste in appraisal records or ask AI to rate someone, and the judgement must be yours (clauses 4 and 7.1(a)).

"Can I use AI to shortlist CVs?" No, not without an approved use and a designed human review. CVs are red data, and shortlisting is a decision about people. Speak to HR.

"Someone recorded our team meeting with an AI note-taker without asking." Remind them of clause 5.2. If the meeting covered HR matters or anything sensitive, ask them to delete the recording and transcript, and report it as an incident if they will not.

"A client has asked whether we use AI on their work. What do I say?" Tell the truth, briefly: which tools, for what, with what safeguards. Check their contract first, and involve [contracts owner] if there is any restriction.

"Do I need to tell clients every time AI helps with an email?" No. Clause 9.2 applies when AI has materially produced work they are paying for, when the contract requires it, or when they would reasonably expect to know.

"Is the AI summary of a document good enough to rely on?" For a quick orientation, often. For anything going to a client, into a number, or into a decision about someone, check it against the source (clause 8).

"Can I connect the AI tool to my inbox or our shared drive?" Only if IT has approved that connection. Some approved tools already have approved connections listed in the register.

"A team member has built a useful automation with AI. What now?" Good. It needs a named owner, a short description and a person who reviews exceptions (clause 8.4). If it acts on systems or sends things externally, it is an agent and goes through clause 6 and the request form.

"Someone pasted payroll data into a free chatbot. What do I do?" Ask them to delete the conversation if they can, keep a screenshot, and report it within 24 hours using Appendix C. The data protection lead will decide whether it is reportable to the ICO. Do not start a disciplinary conversation before the report is made.

"The CEO rang asking me to make an urgent payment to a new account." Call back on the number we already hold for the CEO, whatever the call sounded like. That is the rule for everyone, including the CEO (clause 12.2).

"Can I use AI to translate a document?" Yes, within the colour rules. Have a fluent speaker check anything external or legal; machine translation can change meaning without looking wrong.

"Can I use AI-generated images in marketing?" Yes, in approved tools. Check them for anything that looks like a real person, a brand or a copyrighted character, and label realistic images that could be mistaken for real (clause 9.4).

"Do I have to do the training if I barely use AI?" If you have access to an approved tool, yes. It is short, and it covers the data rules, which apply even if you only use AI occasionally.

"Are we monitoring what people type into AI tools?" We can see the administrative logs that approved tools keep, as explained in the staff privacy notice, and use them to check the policy is followed. We do not use them to assess individual performance without telling people first (clause 15.4).

"Someone keeps breaking the rules. What happens?" Repeated or deliberate breaches go through the normal disciplinary procedure (clause 15.3). A first, honest mistake that is reported promptly is a training issue.

"Who do I ask if the FAQ does not cover it?" [Policy owner], at [contact]. If you had the question, others will too; send it so this FAQ can be updated.

---

12. Sector add-ons

These clauses go into Appendix E of the Word template. Insert only those that apply and renumber if you move them into the main policy. Each sector also has its own regulator guidance, linked, which takes precedence.

Financial services (FCA-regulated firms)

The FCA has said it does not plan to introduce extra regulations for AI, relying instead on existing rules, particularly the Consumer Duty and the Senior Managers and Certification Regime. That means an AI failure is judged against the rules you already follow: good customer outcomes and fair value under the Consumer Duty, accountability of a named senior manager, systems and controls under SYSC (including outsourcing in SYSC 8 and operational resilience in SYSC 15A), and record-keeping of relevant communications under SYSC 10A. Banks, building societies and PRA-designated investment firms should also apply the PRA's model risk management principles (SS1/23) to AI models. Inside information under the UK Market Abuse Regulation is red data.

FS.1 [Name of senior manager, SMF function] is accountable for the firm's use of AI and receives a report on AI use, incidents and customer outcomes at least [quarterly].

FS.2 AI must not be used in any customer communication, product, pricing or decision unless the use has been assessed for its effect on customer outcomes under the Consumer Duty, including customers with characteristics of vulnerability.

FS.3 AI note-takers and transcription tools must not be used in conversations that must be recorded under SYSC 10A unless the tool's records meet our recording and retention requirements and Compliance has approved it.

FS.4 Inside information, and any information that could be inside information, is red data and must not be entered into any AI tool.

FS.5 AI models used in credit, underwriting, pricing, fraud or financial reporting are models for the purpose of our model risk policy and must be inventoried, validated and monitored under it.

FS.6 Material AI tools provided by third parties are assessed under our outsourcing and third-party risk policy before use.

Legal services (SRA-regulated firms)

The SRA published a warning notice on the misuse of AI in August 2026, which it will have regard to in exercising its regulatory functions. Its accompanying news release said it had received 42 reports of potential AI misuse in the twelve months to July 2026. The notice focuses on inaccurate AI output, including fabricated citations of the kind the Divisional Court addressed in R (Ayinde) v London Borough of Haringey [2025] EWHC 1383 (Admin), on client confidentiality when information goes into AI tools without adequate safeguards, and on supervision. Responsibility for the work stays with the solicitor and the firm, whatever tool produced it. The Law Society's guidance, Generative AI: the essentials, is a useful companion.

LG.1 Client information, including the fact that a person is a client, is red data unless the tool is approved for client work in Appendix A and the client's terms of engagement do not prohibit its use.

LG.2 Every citation, quotation and statement of law produced with AI must be checked against an authoritative source by the fee earner before it is relied on, sent to a client or filed with a court or tribunal.

LG.3 Supervisors are responsible for checking that people they supervise follow LG.2, and for reviewing AI-assisted work to the same standard as any other work.

LG.4 Privileged material must not be entered into any AI tool that is not approved for privileged material in Appendix A.

LG.5 Our standard terms of business tell clients how we use AI on their matters. Where a client has asked us not to use AI, the matter is flagged and LG.1 applies to all of it.

Healthcare and care providers

Health data is special category data under UK GDPR, so most clinical use is red by default. NHS England's guidance on AI-enabled ambient scribing products expects organisations adopting them to complete a DPIA, appoint a clinical safety officer and complete clinical risk management documentation under the DCB0160 standard, and to consider whether a product qualifies as a medical device regulated by the MHRA. The Caldicott principles and the duty of confidentiality apply alongside UK GDPR. The guidance applies in England; Scotland, Wales and Northern Ireland have their own arrangements.

HC.1 Patient and service user information is red data. It may only be entered into AI tools approved for clinical use in Appendix A, each supported by a DPIA and, where it supports care, a clinical safety case under DCB0160 with a named clinical safety officer.

HC.2 AI must not be used to diagnose, triage, prescribe or make treatment decisions unless the product is appropriately regulated as a medical device for that purpose and its use has been approved through clinical governance.

HC.3 Ambient scribing and transcription tools may only be used with the patient's knowledge, and clinicians must review and correct every AI-generated note before it is saved to the record.

HC.4 Clinicians remain professionally accountable for any AI-assisted record, letter or decision.

HC.5 Clinical incidents involving AI are reported through our clinical incident system as well as under clause 14.

Education (schools, colleges and education providers)

The Department for Education's policy paper on generative AI in education sets out its position, and its generative AI product safety expectations (updated as standards in January 2026) describe what safe AI products for schools should do, including filtering, monitoring and data protection. Keeping children safe in education links to them from its filtering and monitoring section. For assessed work, the JCQ's AI Use in Assessments guidance (updated April 2025) treats undeclared AI use as malpractice. Children's data merits specific protection under UK GDPR, and the Data (Use and Access) Act added a duty to consider children's higher protection matters when designing online services.

ED.1 Pupil and student personal data is red data. Staff may only use AI tools approved for pupil data in Appendix A, each supported by a DPIA.

ED.2 Pupils may only use AI tools that [the school] has approved for pupil use, configured with appropriate filtering and monitoring, and in line with our acceptable use agreement for pupils.

ED.3 Teachers remain responsible for marking, feedback and reports. AI may assist, but a teacher reviews every grade and comment before it is shared.

ED.4 Use of AI in assessed or coursework tasks follows the current JCQ guidance and our malpractice policy. Pupils must be told what AI use is allowed for each task and how to acknowledge it.

ED.5 Safeguarding concerns that arise through AI tools, including harmful content or disclosures, are reported to the designated safeguarding lead immediately.

Recruitment and HR (any employer, and recruitment businesses)

Recruitment is where AI meets the most law at once: UK GDPR, the new Articles 22A to 22D, and the Equality Act. The ICO's Recruitment rewired report (March 2026) found many employers lacked meaningful human involvement and consistent safeguards, and expects transparency to candidates, consistent human involvement and bias monitoring. DSIT's Responsible AI in Recruitment guide (2024) sets out assurance questions to ask suppliers. If candidates are in the EU, AI used in recruitment is high-risk under Annex III of the EU AI Act, with deployer obligations now due from 2 December 2027.

RC.1 AI tools used to source, screen, rank, assess or interview candidates, or to assess staff, must be approved under clause 11, supported by a DPIA, and listed in Appendix A.

RC.2 Every decision to reject, shortlist, hire, promote or dismiss is made by a trained person who reviews the underlying evidence, can depart from the AI's recommendation, and records the reason for the decision.

RC.3 Candidates are told in the job advert or application process when AI is used, what it does, and how to ask for a human review or a reasonable adjustment.

RC.4 Before use and at least every [six months], [HR] checks outcomes by protected characteristic where data allows, and asks the supplier for its bias testing evidence.

RC.5 AI tools must not be used to infer protected characteristics, emotions or personality from video, voice or writing.

RC.6 Recruitment businesses acting for clients agree in writing with each client who is responsible for each step of any AI-assisted process.

13. Worked example

Illustrative example. Calder Freight Ltd, a fictional 180-person logistics business in West Yorkshire, adopts the policy in October. Everything below is invented to show the method.

Starting point. The anonymous survey (62% response) shows 58% of staff using AI for work, almost all on personal accounts. Finance uses a free chatbot to tidy month-end commentary; sales uses a meeting note-taker that joins every customer call automatically; two developers use a coding assistant on personal licences; HR has trialled a CV ranking feature in its applicant tracking system.

Decisions in the drafting week.

ISSUE	DECISION	CLAUSE
No approved tool	Approve Copilot for all staff (amber), already within the Microsoft licence; approve a business coding assistant for the two developers (amber, no secrets)	3, Appendix A
Finance commentary	Move to Copilot; management accounts are amber; unreleased results remain red until published	4
Sales note-taker	Withdraw the unapproved tool; enable the approved platform's note-taker for internal meetings only; customer calls only with agreement	5
CV ranking	Switch off pending a DPIA and a designed human review; HR to complete Appendix B	7.1(a), 8.5, RC.1 to RC.5
EU exposure	Calder runs a depot in Rotterdam with 14 staff: Article 4 applies to that entity; training record to include them	13
Payments	Accounts payable already calls back on bank detail changes; extend to all urgent requests from senior staff	12.2

Rollout. The CEO announces the policy and the amnesty. Eleven tools are disclosed in 30 days; three are approved through the request form, eight are withdrawn. Training runs over three weeks in groups of ten, each person working on a task from their own job. At week six, the repeat survey shows personal-account use for work down to 19% of staff, mostly for green tasks. One incident is reported: a transport planner pasted a driver's sickness note into Copilot. It was amber-approved, not red-approved; the data protection lead assesses it as a low-risk breach that does not need reporting to the ICO, records the reasons, and the planner's team gets a ten-minute refresher on health data.

First review at month six. Two changes: the note-taker is approved for customer calls where the customer agrees at the start, and a request to use an AI agent to draft replies to delivery queries is approved with clause 6.3 approvals in place for any reply mentioning compensation.

14. Common mistakes

Banning everything. A blanket ban does not stop use; it moves it to personal phones, where you have no terms, no logs and no say. Approve a business tool and set rules for it.

Writing the policy before choosing the tool. The rules depend on the tool's terms. Approving a tool after the policy means rewriting the data rules.

Vague data rules. "Do not input confidential information" leaves every employee to decide what confidential means. Use the traffic-light table and fill it with your own examples.

Rubber-stamp human review. "A human will review" is not enough when the reviewer approves every recommendation without seeing the evidence. Define what the review involves (clause 8.5).

Forgetting AI inside existing software. Vendors switch AI features on in CRM, HR, finance and video software with routine updates. Ask IT to check settings quarterly and put each feature in the register.

Ignoring note-takers. They join calls automatically and send summaries to all attendees, including external ones. Decide the rules before someone records a disciplinary hearing.

Letting agents borrow personal credentials. An agent running as you can do anything you can, and its actions look like yours in the logs.

No request route, or a slow one. If approval takes months, people will not wait. Set a target of ten working days and publish decisions.

Treating the policy as the training. Asking staff to read and sign a policy produces signatures, not behaviour. Teach the rules on real work.

No owner, no date. Provider terms and ICO guidance are changing. A policy without a named owner and review date is out of date within months.

Copying a US template. US templates refer to state laws, at-will employment and US privacy regimes, and miss UK GDPR, the Equality Act and the ICO. Several popular free templates still do.

Over-reading the EU AI Act. A UK-only business does not need an AI Act compliance programme, and an EU-facing one needs advice on its specific position, not a template clause. Article 4 asks for measures and a record, not certificates.

15. Glossary

TERM	MEANING IN THIS PACK
AI agent	An AI tool that can take actions on systems, such as sending emails, updating records or making payments
AI literacy	The skills and understanding to use AI tools well and recognise their risks; the subject of Article 4 of the EU AI Act
Amber data	Internal, non-sensitive business data; approved business accounts only

TERM	MEANING IN THIS PACK
Automated decision-making (ADM)	A decision taken by automated means; "solely automated" under UK GDPR Article 22A if there is no meaningful human involvement
Business account	An account under a business or enterprise plan, with contract terms the business has reviewed
Controller	The organisation that decides why and how personal data is processed
Deployer	Under the EU AI Act, an organisation using an AI system in a professional capacity
DPIA	Data protection impact assessment, required under UK GDPR Article 35 for processing likely to result in high risk
Green data	Public or anonymised information; any approved tool
Hallucination	Confident, plausible output that is wrong, such as an invented figure or citation
Meaningful human involvement	Review by someone with the competence and authority to change the outcome, who considers the evidence before the decision takes effect
Processor	An organisation that processes personal data on behalf of a controller, such as an AI provider under a business plan
Prompt injection	Instructions hidden in content an AI tool reads (an email, a web page, a file) that cause it to act against the user's intent
Provider	Under the EU AI Act, an organisation that develops an AI system and places it on the market under its own name
Red data	Sensitive data that goes into no AI tool without specific written approval
Shadow AI	AI tools used for work without the organisation's approval
Special category data	Health, ethnicity, religion, sexual orientation, trade union membership, genetic and biometric data, and similar, under UK GDPR Article 9

16. Sources and further reading

Legislation and regulators

- UK GDPR and Data Protection Act 2018: ICO guidance index, Information Commissioner's Office.
- Data (Use and Access) Act 2025, section 80 and section 103, legislation.gov.uk; commencement in SI 2026/82.
- Guidance on AI and data protection, ICO (under review).
- Consultation on draft guidance on automated decision-making and profiling, ICO, 2026.
- Recruitment rewired, ICO, 2026.
- Monitoring workers, ICO.
- Personal data breaches: a guide, ICO.
- ICO intervention into AI recruitment tools, ICO, 2024.
- Equality Act 2010 and Digital Markets, Competition and Consumers Act 2024, legislation.gov.uk.

- Report on Copyright and Artificial Intelligence, DSIT and DCMS, March 2026.
- Getty Images v Stability AI, judgment, High Court, November 2025.
- EU AI Act, Regulation (EU) 2024/1689 and Digital Omnibus on AI, Regulation (EU) 2026/1744, EUR-Lex.
- AI literacy questions and answers, European Commission.
- AI Act obligations in force and delays under omnibus 2026/1744, ActuIA, 2026 (secondary summary of dates).
- Our approach to AI, Financial Conduct Authority.
- Misuse of AI warning notice, Solicitors Regulation Authority, August 2026.
- Guidance on the use of AI-enabled ambient scribing products, NHS England.
- Generative AI in education, Department for Education.
- Responsible AI in Recruitment, DSIT, 2024.
- Artificial Intelligence Playbook for the UK Government, DSIT, 2025.

Statistics cited

- Rise in shadow AI tools raising security concerns for UK, Microsoft UK and Censuswide, October 2025: 71% of 2,003 UK employees had used unapproved consumer AI tools at work; 28% said their employer offered no approved option.
- Artificial intelligence in UK businesses: 2023 to 2026, ONS, July 2026: around 35% of businesses with ten or more employees use AI; 11% say more than half their workforce has had AI-related training.
- Arup deepfake scam, CNN, May 2024.

Other free templates worth reading

We reviewed the best free AI policy templates and guides available to UK businesses. Each has something to offer; this pack aims to combine their strengths with the forms and sector material none of them provides in full.

RESOURCE	WHAT IT DOES WELL	WHERE IT STOPS
Spotdev AI policy template	Detailed sixteen-clause UK policy with reasoning per clause; strong on agents, security and shadow AI; tools and agent registers	No request form, incident form or staff summary; one bracketed sector clause
CIPD guide to preparing your organisation for AI use	Good on the people side: cross-functional ownership, communication, HR uses	Policy design checklist is for members; not a ready-to-use policy
Fisher Phillips AI policy templates	Family of templates including note-takers and vendor management	Written for US employers and US law
Rocket Lawyer UK AI policy checklist	Clear on non-contractual status and publishing to staff	A post-drafting checklist, not the policy text
Socitm template policy for generative AI	Public sector template aligned to ICO expectations	Written for councils and public bodies
UK Government AI Playbook	Ten clear principles; good on understanding limits and human control	Guidance for government, not a staff policy
Sprintlaw guide to generative AI policies	Good summary of UK legal pressure points	Guide rather than template

RESOURCE	WHAT IT DOES WELL	WHERE IT STOPS
Thirdfort AI policy template for law firms	Focused on legal practice risks	Legal sector only

The companion article, AI policy template for UK businesses, explains how to write the policy in a week and why training matters more than the document.

17. About this resource

Written by Anyway Ventures, the York technology and AI practice led by Brannan Coady, drawing on putting AI tools and policies into UK businesses. Version 1.0, October 2026; law and guidance checked to 25 October 2026. Free to use, adapt and share with attribution. Not legal, tax or investment advice.

ABOUT ANYWAY VENTURES

Senior hands for technology and AI decisions

Technology and AI work for private equity funds, venture investors and UK mid-market management teams: before a deal, during the hold and at exit. Every engagement is led by Brannan Coady, an exited founder and former CEO. Fixed price, built to be handed over.

FOR INVESTORS

Technical due diligence

Code, architecture, security, IP and one-to-one team interviews, every material finding costed. Red-flag review in five to seven days.

ANYWAYVENTURES.COM/TECHNICAL-DUE-DILIGENCE

AI due diligence

Whether a target's AI claims are real, defensible and theirs, and a costed AI value-creation plan for the IC.

ANYWAYVENTURES.COM/AI-DUE-DILIGENCE

AI disruption risk review

Is AI a threat to the business or its moat? One deal, or every holding in a portfolio ranked by exposure.

ANYWAYVENTURES.COM/AI-DISRUPTION-RISK

Technical due diligence for VCs

A two to three day check before the term sheet: the product tested on evidence, not the demo.

ANYWAYVENTURES.COM/VC-TECHNICAL-DUE-DILIGENCE

Retained AI operating partner

Senior AI capacity held by the fund and deployed into the company with the strongest AI value case.

ANYWAYVENTURES.COM/AI-OPERATING-PARTNER

Exit readiness and vendor technical due diligence

Find what a buyer's diligence will find, 12 to 24 months out, fix it, then publish a report buyers can rely on.

ANYWAYVENTURES.COM/EXIT-READINESS

FOR OPERATORS

AI audit and readiness assessment

Two to three weeks to a ranked, costed list of AI opportunities, the gaps to close and a 90-day plan.

ANYWAYVENTURES.COM/AI-AUDIT

Embedded AI delivery

Internal tooling, automation and AI agents in production inside your business, built with your team.

ANYWAYVENTURES.COM/EMBEDDED-AI-DELIVERY

Fractional Chief AI Officer

One senior person who owns AI on your leadership team, two to four days a month.

ANYWAYVENTURES.COM/FRACTIONAL-CHIEF-AI-OFFICER

Hands-on AI training

In-person training where every attendee builds a working tool on their own job.

ANYWAYVENTURES.COM/AI-TRAINING

Board and NED roles

Non-executive seats where technology, product and AI decide the outcome.

ANYWAYVENTURES.COM/BOARD-AND-NED

Every enquiry starts with the same free 30-minute call.

A straight answer on whether we can help, what it would cost and how long it would take.

EMAIL
brannan@anywayventures.com

BOOK A CALL
anywayventures.com/contact

MORE RESOURCES
anywayventures.com/resources

Anyway Ventures Limited · Registered in England and Wales, company number 16318119 · Registered office: Newbury House, Main Street, Little Ouseburn, York, YO26 9TD. This resource is general guidance, not legal, tax or investment advice. Free to use, adapt and share with attribution to Anyway Ventures. Version 1.0, October 2026.