Artificial Intelligence Usage Policy

Template version 1.0, updated July 27, 2026. Free to copy, edit, and adopt under your own name. Published by Velum, https://velumprivacy.com. This is a template, not legal advice.

Document owner: [ROLE, e.g. Data Protection Officer]

Approved by: [APPROVING BODY]

Effective date: [DATE]

Version: 1.0

Next review: [DATE, suggested 12 months]

> This is a template, not legal advice. It is written to be a sound starting point across sectors and jurisdictions. Have your data protection officer or counsel review it before adoption. Three clauses are marked **[CHECK LOCAL LAW]** because the correct answer genuinely depends on where you operate.

---

1. Purpose and scope

This policy sets out how artificial intelligence tools may be used at [ORGANISATION], what may be entered into them, and what may not.

It applies to all employees, contractors, temporary staff, interns, and volunteers, on any device, whether or not the device is owned by [ORGANISATION], when the work being done is [ORGANISATION] work.

It covers any tool that sends text, files, images, audio, or code to a system that generates or analyses content. That includes general assistants, coding assistants, meeting transcription services, translation tools, and AI features built into software [ORGANISATION] already licenses.

The last category is the one policies usually miss. An AI summarisation feature inside an existing product is in scope even though nobody ever made a decision to adopt an AI tool.

---

2. Definitions

AI tool. Any software that processes an input to generate content, a summary, a classification, a translation, or code, using a machine learning model. Includes standalone products and features embedded in other products.

Public AI tool. An AI tool where inputs are transmitted to a provider outside [ORGANISATION]'s control, under that provider's own terms.

Prompt. Anything submitted to an AI tool: typed text, pasted text, uploaded files, images, screen shares, audio, or code.

Output. Anything an AI tool returns.

Confidential information. As defined in [EXISTING CONFIDENTIALITY OR INFORMATION SECURITY POLICY]. This policy does not create a second definition.

Personal data. Any information relating to an identified or identifiable person, as defined by [APPLICABLE LAW, e.g. UK GDPR, EU GDPR, or your local equivalent].

---

3. Roles and responsibilities

[ROLE, e.g. the Data Protection Officer or Head of Information Security] owns this policy, maintains the approved tool register in Annex A, and decides exception requests.

Line managers are responsible for their team knowing that this policy applies to them, and for raising a request rather than allowing a quiet workaround.

All staff are responsible for what they put into a prompt and for what they do with an output. Using an AI tool does not transfer responsibility for the work to the tool.

[ROLE, e.g. IT] maintains the technical controls that support this policy and reports on tool usage to the policy owner [FREQUENCY].

Exception requests go to [ROLE OR MAILBOX]. A decision is given within [NUMBER] working days.

---

4. Approved tools

Only tools listed in Annex A may be used for [ORGANISATION] work.

Tools are listed in one of three tiers.

Tier 1, approved. Use freely, within the data rules in section 5. These are tools where [ORGANISATION] holds a contract with terms we have reviewed, covering training, retention, and where data is processed.

Tier 2, conditional. Use only for the purposes and data tiers listed against the tool in Annex A. These are tools with an acceptable use case and a limitation, for example a provider that processes data outside [JURISDICTION], or a free tier whose terms permit training on inputs.

Tier 3, prohibited. Do not use for any [ORGANISATION] work. This tier exists so that specific tools can be named, because "not approved" and "banned" are read differently.

A tool that is not in Annex A is not approved. Section 12 explains how to ask for one to be added. Do not assume that a tool being widely used elsewhere in the industry means it has been assessed here.

Personal accounts on approved tools are not approved. Use the account [ORGANISATION] provides. A personal account puts the work outside our contract, outside our retention terms, and outside our ability to retrieve or delete anything.

---

5. What may be entered into an AI tool

This is the operative clause of the policy. Everything else supports it.

Data is classified in the tiers defined in [EXISTING DATA CLASSIFICATION POLICY], or in Annex B if [ORGANISATION] has no existing scheme.

Data tier	Public AI tools	Tier 1 approved tools	What this covers
Public	Permitted	Permitted	Already published or approved for publication
Internal	Not permitted	Permitted	Ordinary business information, not sensitive
Confidential	Not permitted	Permitted with masking or written approval	Client work, commercial terms, unpublished plans, source code
Restricted	Never	Never	Personal data of clients, staff, or patients. Special category data. Regulated records. Credentials. Privileged material

Adjust these tiers to match the classification scheme [ORGANISATION] already uses rather than introducing a parallel one. Two schemes means neither gets used.

The rule staff need to remember is short: **if the tool does not need to know who the person is in order to do the task, the person's identity does not belong in the prompt.**

---

6. Prohibited uses

The following are prohibited regardless of tool tier.

1. Entering personal data of clients, customers, patients, or staff into any AI tool not approved for that data tier.
2. Entering special category data, including health, biometric, genetic, racial or ethnic origin, political opinions, religious beliefs, trade union membership, sex life, or sexual orientation data, into any public AI tool.
3. Entering credentials, API keys, access tokens, or security configuration into any AI tool.
4. Entering material covered by legal professional privilege into any tool not explicitly approved for it.
5. Entering third-party confidential information where our contract with that third party does not permit disclosure to a subprocessor. See section 9.
6. Using AI output as the sole basis for a decision that materially affects a person, including decisions on employment, credit, benefits, admission, discipline, or care. See section 7.
7. Presenting AI-generated work as though it were produced without AI where a client, a court, a regulator, or a publisher requires disclosure.
8. Using AI tools to generate content that would breach [ORGANISATION]'s existing policies on conduct, discrimination, or harassment if a person had written it.
9. Circumventing a technical control that implements this policy.

7. Human review and accountability for output

A person is accountable for every output used in [ORGANISATION] work. The tool is not.

Before output is sent to a client, filed, published, or committed, the person using it must:

- Verify every factual claim, figure, citation, and legal reference against a source. AI tools produce confident, well-formatted, wrong answers, and citations to material that does not exist are a known and recurring failure.
- Confirm the output does not reproduce someone else's protected material.
- Confirm the output is appropriate for the audience and consistent with [ORGANISATION]'s standards.

For decisions with a material effect on a person, a qualified human makes the decision. AI output may inform it. **[CHECK LOCAL LAW: rules on automated decision-making vary, and requirements covering AI in employment, credit, and healthcare decisions are changing quickly.]**

8. Transparency and disclosure

[ORGANISATION] discloses its use of AI where a client, regulator, court, publisher, or counterparty requires it, and where a reasonable person would expect to be told.

Disclose when:

- A client contract or a professional body rule requires it.
- Output is submitted to a court, a regulator, or a public body with rules on AI-assisted submissions.
- A person is interacting with an AI system and might reasonably believe they are interacting with a person.
- [ADD YOUR OWN CATEGORIES]

Routine internal assistance with drafting, summarising, and formatting does not require disclosure unless the above applies.

9. Client and third-party data

Where [ORGANISATION] holds data belonging to a client or another organisation, our contract with them governs what we may do with it. That contract usually predates AI tools and rarely mentions them.

Before entering third-party data into any AI tool, confirm that:

1. The contract permits disclosure to subprocessors, and the AI provider fits within that permission.
2. Any subprocessor notification or approval requirement has been met.
3. The client has not issued a specific instruction about AI use. Some have. Check.

Where the position is unclear, treat the data as Restricted until [ROLE] confirms otherwise.

In most cases the practical alternative is to remove the identifying and confidential content before the prompt leaves, so that no third-party data is disclosed at all. Section 16 covers this.

---

10. Data protection

Entering personal data into an AI tool is a processing operation and requires a lawful basis. Where the provider processes data outside [JURISDICTION], it may also be a restricted international transfer requiring safeguards. **[CHECK LOCAL LAW: transfer rules and adequacy findings differ by jurisdiction and change over time.]**

Accordingly:

- A new AI tool that will process personal data requires a data protection impact assessment before approval, or a documented assessment that one is not required.
- The record of processing activities is updated when an AI tool is approved for personal data.
- Data subject rights, including access, rectification, and erasure, apply to data held by an AI provider on our behalf. Confirm before approval that we can meet those requests.
- Data minimisation applies to prompts exactly as it applies everywhere else. A prompt carrying a whole record in order to answer a narrow question is not minimised.

---

11. Intellectual property and confidentiality

Output from AI tools may not be protectable by copyright in every jurisdiction, and may resemble material the model was trained on. Consult [ROLE] before relying on output as [ORGANISATION] intellectual property.

Inputs to a public AI tool may be retained by the provider and may be used to improve its models, depending on that provider's terms. Entering unpublished confidential material into such a tool risks the confidentiality of that material and may breach an obligation we owe to someone else.

Source code is confidential information. Code containing credentials, proprietary algorithms, or client-specific logic follows the same rules as any other confidential material.

---

12. Requesting a new tool

To request that a tool be added to Annex A, send [ROLE OR MAILBOX]:

1. The tool, the vendor, and the plan or tier.
2. What you want to use it for, specifically.
3. What data would go into it.
4. Whether a free tier or a personal account is involved.

[ROLE] assesses the request against the provider's terms on training and retention, where processing happens, security posture and certifications, the subprocessor list, deletion and export capability, the contractual terms available to us, and whether an approved tool already covers the need.

Do not start using a tool while a request is pending.

---

13. Records and monitoring

[ORGANISATION] [DOES / DOES NOT] monitor AI tool usage on managed devices and networks. Where monitoring occurs it is carried out under [EXISTING MONITORING POLICY] and staff are informed of it. **[CHECK LOCAL LAW: workplace monitoring rules and works council consultation requirements vary significantly, and getting this wrong is itself a compliance failure.]**

[ROLE] maintains the approved tool register and a record of exception decisions.

---

14. Incidents

If confidential or personal data has been entered into an AI tool in breach of this policy, report it to [ROLE OR MAILBOX] within [NUMBER] hours of realising.

Report it even if you are not sure. Report it even if it was you. Do not delete the conversation before reporting, because the record may be needed for the assessment.

Handling follows [EXISTING INCIDENT RESPONSE PROCEDURE]. The assessment considers what was disclosed, to whom, under what terms, whether it can be retrieved or deleted, and whether the event is notifiable to a regulator or to affected individuals under [APPLICABLE LAW].

Reporting in good faith will not result in disciplinary action. Concealing an incident may. This wording is deliberate. A policy people are afraid of produces silence, and silence turns a small disclosure into an undetected one.

---

15. Training

All staff complete AI usage training at induction and [FREQUENCY] thereafter. Staff handling [REGULATED DATA CATEGORY] complete the additional module covering [SPECIFICS].

Training covers what this policy permits, how to recognise data that must not be entered, how to check output, and how to report an incident.

---

16. Reducing exposure at the boundary

A policy controls behaviour only as far as people follow it, and prompt-by-prompt judgement is the weakest control in any information security programme. Where possible, [ORGANISATION] prefers controls that make a breach of section 5 difficult rather than merely forbidden.

That includes:

- Removing or masking personal and confidential data from text before it reaches an AI tool.
- Providing approved tools good enough that workarounds are not tempting.
- Making the approved path faster than the unapproved one.

[DESCRIBE THE CONTROLS YOU ACTUALLY HAVE. Delete this section if you have none, rather than describing an intention as though it were a control.]

17. Review

This policy is reviewed [FREQUENCY] by [ROLE], and additionally when a significant new tool is adopted, when the regulatory position changes, or after any incident under section 14.

Annex A: Approved tool register

Tool	Vendor	Tier	Approved for	Data tiers permitted	Account type	Reviewed
[Tool]	[Vendor]	1	[Purposes]	Public, Internal, Confidential	[ORGANISATION] SSO	[DATE]
[Tool]	[Vendor]	2	[Named purposes only]	Public, Internal	[ORGANISATION] SSO	[DATE]
[Tool]	[Vendor]	3	Prohibited	None	None	[DATE]

Keep this register somewhere staff can find in under ten seconds. A register nobody can locate is a register nobody consults.

Annex B: Data classification, short version

Use this only if [ORGANISATION] has no existing classification scheme. If one exists, delete this annex and reference the real one.

Public. Published, or approved for publication. Marketing material, public filings, the website.

Internal. Ordinary business information whose disclosure would be unwelcome but not harmful. Internal process documents, routine team communication.

Confidential. Disclosure would cause harm to [ORGANISATION], a client, or a partner. Client deliverables, commercial terms, unpublished plans, source code, security documentation.

Restricted. Disclosure would cause serious harm, or the data is regulated. Personal data, special category data, financial records, medical records, privileged material, credentials, and anything covered by a specific confidentiality undertaking.

Annex C: Incident report form

Field	Entry
Reported by	
Date and time of the event	
Date and time noticed	
Tool involved	
Account used (organisation or personal)	
What was entered	
Data tier	
Number of people whose data was involved	
Whose data (clients, staff, patients, third party)	
Conversation deleted or retained	
Any output shared onward	
Immediate action taken	

Send to [ROLE OR MAILBOX]. Do not include the affected data itself in this form. Describe it.

Annex D: One-page card for staff

Circulate this on its own. It is the part people will read.

Using AI at [ORGANISATION]

Use these tools: [LIST]. Sign in with your [ORGANISATION] account, never a personal one.

Never put these into any AI tool:

- Names, contact details, or anything identifying a client, customer, patient, or colleague
- Health, financial, or other regulated records
- Passwords, API keys, access tokens
- Anything a client has told us to keep confidential
- Unpublished contracts, commercial terms, or source code, unless the tool is approved for it

Safe to use it for: drafting, rewriting, summarising your own notes, explaining a concept, generating code that contains nothing specific to a client, formatting, and translating public material.

Before the output goes anywhere: check every fact, figure, name, and citation. The tool will be confidently wrong sometimes, and it is your name on the work.

If you paste something you should not have: tell [ROLE OR MAILBOX] straight away. Do not delete the chat first. Nobody is in trouble for reporting.

Not sure? Ask [ROLE OR MAILBOX] before you paste, not after.