

[COMPANY NAME] AI Use Policy

Version: [1.0] · **Effective date:** [DATE] · **Policy owner:** [NAME, JOB TITLE] · **Next review:** [DATE, no more than 12 months from the effective date]

Guidance for the owner (delete this box before you share the policy). Replace every [BRACKET] with your own details. Where the policy offers a choice, such as [OPTION A / OPTION B], keep the one that fits and delete the other. Filling it in takes about 10 minutes; the rollout guide covers the rest of the 30-minute setup. **Not legal advice. Have your attorney review it if you handle health, legal or financial data,** or if your industry, your contracts or the places where you do business have their own rules about customer data or automated tools.

1. Purpose

[COMPANY NAME] ("we", "the company") lets staff use artificial intelligence (AI) tools, such as chat assistants, writing helpers and AI features built into the software we already pay for, when they help us do good work faster. This policy explains which tools are allowed, what information may and may not go into them, how AI output is checked before anyone relies on it, and what to do when something goes wrong.

The goals are simple:

1. Keep customer, staff and company information private.
2. Make sure a person checks anything AI produces before it is used.
3. Be honest with customers about how we use AI.
4. Give staff clear answers so nobody has to guess.

2. Who this policy covers

This policy applies to everyone who works for or with [COMPANY NAME]: employees (full-time, part-time and seasonal), owners and managers, [contractors and freelancers / interns / volunteers] who handle company or customer information. It covers any AI tool used for company work, on any device, including personal phones and home computers, and including AI features inside apps such as email, word processing, spreadsheets, design, accounting and customer messaging tools.

3. Approved tools

Staff may use **only** the AI tools listed in the company's **Approved AI Tools Register** for company work. The register lists each tool, the plan or account type we use, who approved it, what kind of data may go into it, and when it was last reviewed.

- The current register is kept at: [LOCATION, for example "the shared drive, Policies folder" or "pinned in the staff chat"].
- **Use the company account, not a personal one.** If the register says we use a business or team plan, sign in with the company login for that plan. A personal free account of the same tool may handle data differently.
- If a tool isn't on the register, it isn't approved, even if it's free, popular or built into another app. Ask first (section 6).

Guidance for the owner. AI companies often treat business and personal accounts differently. For example, several major providers don't use content from their business, team or enterprise plans to train their AI models by default, while some personal and free plans do use chats for training unless the user turns that off. Check each tool's own privacy or data-controls page before you approve it, and write what you find in the register's notes column. The settings checklist at the end of this policy lists the settings to look for.

4. Never paste these into any AI tool

Unless a specific tool is approved for that data type in the register **and** the owner has approved it in writing, staff must never type, paste, upload, dictate or screenshot the following into an AI tool:

1. **Passwords, PINs, security codes, backup codes or login links**, for any account.
2. **Payment card numbers**, expiration dates and security codes.
3. **Bank account and routing numbers**, and wire or payment instructions.
4. **Social Security numbers**, tax ID numbers, driver's license, passport or other government ID numbers.
5. **Health information** about any customer or staff member, including medical notes, conditions, prescriptions, injuries and insurance claims.
6. **Customer records with names, addresses, phone numbers or email addresses**, such as customer lists, contact exports and order histories.
7. **Staff personnel information**: pay, reviews, discipline, medical leave, background checks and home addresses.
8. **Legal matters**: contracts under negotiation, legal advice from our attorney, claims, disputes and anything marked confidential.
9. **Financial records** that aren't public: the company's accounts, payroll, tax returns, pricing agreements and forecasts.
10. **Information a customer or partner gave us in confidence**, or that a contract or nondisclosure agreement (NDA) says we must protect.
11. **Security details**: Wi-Fi passwords, alarm codes, door codes, network settings, system access keys and API keys.
12. **Photos or recordings of people** (customers, staff, children) taken for work, unless the person agreed and the use is approved.

[ADD ANY ITEMS SPECIFIC TO YOUR BUSINESS, for example "patient intake forms", "client case files" or "student records".]

When in doubt, leave it out. If you're not sure whether something is on this list, don't paste it. Ask [NAME OR ROLE].

5. What is OK to paste

These are generally fine to use with an approved tool:

- Information that's **already public**, such as our website text, published prices, public reviews (with reviewer names removed) and our social posts.
- **Your own drafts** of marketing copy, social posts, product descriptions, job ads and internal how-to notes.

- **General questions** that don't mention a customer or staff member, for example "How do I write a polite reminder for an overdue invoice?"
- **Information with personal details removed.** Replace names with roles ("the customer", "Employee A"), and delete addresses, phone numbers, emails, account numbers and anything else that could identify a person.
- **Spreadsheet formulas, templates and formatting help**, using made-up sample data.

The "never paste" list in section 4 always wins over this list.

6. Who approves new tools

- Only **[NAME, JOB TITLE]** (or, in their absence, **[BACKUP NAME, JOB TITLE]**) can add a tool to the register.
- To request a new tool, send **[NAME]** the tool's name and website, what you want to use it for, what kind of data would go into it, and whether it's free or paid.
- Before approving, the approver checks the tool's own privacy and data pages for: whether our content is used to train AI models and whether that can be turned off; how long chats and files are kept and whether they can be deleted; whether people at the company that makes the tool may review content; and whether a business or team plan is available with admin controls.
- Approval is recorded in the register with the date, the plan we use, the approver's name and the data allowed.
- **AI features inside existing apps count as new tools.** If an app we already use adds an AI assistant, don't use it for customer or staff information until it's on the register.

7. Checking AI output

AI tools can be wrong, out of date, or confidently make things up. **A person is responsible for everything we use, send or publish, whether or not AI helped write it.**

Before using AI output, staff must:

1. **Read all of it.** Don't send or publish anything you haven't read in full.
2. **Check facts, numbers, dates, prices, names and quotes** against a trusted source (our own records, the original document, an official website).
3. **Check that it matches our policies**, such as prices, refund terms, warranties and promises to customers. AI must never invent a discount, a guarantee or a deadline.
4. **Never rely on AI alone** for legal, medical, tax, safety or financial advice to anyone, or for decisions about hiring, firing, pay, credit or pricing for a particular person. A qualified person makes those decisions.
5. **Check images and designs** for wrong text, odd details and anything that looks like someone else's logo, artwork or likeness.
6. **Run code or formulas** on a copy first, never on live company data.

8. Labeling AI-written customer messages

[CHOOSE ONE AND DELETE THE OTHER]

[OPTION A: Label automated replies] Any message a customer receives that was **sent automatically by an AI tool without a person reviewing it first** (for example, a website chat assistant or an automatic email reply) must say clearly that it comes from an automated assistant, and tell the customer how to

reach a person. Messages a staff member drafted with AI help, then read, edited and sent themselves, don't need a label.

[OPTION B: Label all AI-assisted messages] Any customer message where AI wrote a substantial part of the text must include this line: "[This message was drafted with help from an AI tool and reviewed by our team.]"

In every case:

- **Never let an AI tool pretend to be a named person** on our team, and never tell a customer they're talking to a person when they're talking to an automated assistant.
- If a customer asks whether they're talking to a person or AI, tell them the truth.
- Customer-facing chat assistants must hand the conversation to a person when the customer asks.

Guidance for the owner. Some states and industries have rules about disclosing automated chat or AI use to customers. Ask your attorney whether any apply to you, and choose the stricter option if you're unsure.

9. Privacy and data retention

- **Turn off training where you can.** For every approved tool, the approver sets the account so our content isn't used to train the provider's AI models, if the tool offers that setting, and notes it in the register.
- **Delete what you don't need.** Staff should delete chats and uploaded files that contain company information once the task is finished, unless [NAME] asks for them to be kept. [Use temporary or private chat modes where the tool offers them.]
- **Know that deletion may not be instant.** Some providers keep chats for a period after you delete them (for example, for safety or legal reasons), and content reviewed by people at the provider may be kept longer. Record each provider's retention rules in the register's notes.
- **Don't connect AI tools to company accounts** (email, cloud drives, calendars, customer systems) unless the connection is approved in the register.
- **Leaving the company.** When someone leaves, [NAME] removes their access to every company AI account the same day, along with their other accounts.
- **Customer requests.** If a customer asks what we do with their information, or asks us to delete it, tell [NAME OR ROLE]. Don't answer from memory.

10. If something goes wrong: 5 steps

Something has gone wrong if, for example: someone pasted information from the "never paste" list into an AI tool; an AI-written message with wrong information reached a customer; an AI tool or account was used by someone who shouldn't have access; or a tool behaved in a way that exposed company or customer data.

1. **Stop.** Stop using the tool for that task. Don't delete anything yet, so the facts can be checked.
2. **Tell [NAME OR ROLE] right away,** at [PHONE / EMAIL / CHAT CHANNEL]. Say what happened, when, which tool and what information was involved. Nobody will be punished for reporting a mistake promptly. Hiding one is a policy violation.
3. **Contain it.** [NAME] decides on the next steps, such as deleting the chat or file from the tool, changing any password or code that was exposed, signing the account out of all devices, contacting the tool provider, or correcting a wrong message to a customer.

4. **Decide who else needs to know.** If customer or staff personal information may have been exposed, [NAME] talks to [OUR ATTORNEY / INSURANCE PROVIDER] about whether the law or our contracts require us to notify anyone. Breach-notification rules vary from state to state.
5. **Learn from it.** Write down what happened and what we changed. Update this policy, the register or staff training if needed.

11. Breaking the policy

Breaking this policy may lead to [retraining / loss of access to AI tools / disciplinary action up to and including termination, in line with our employee handbook]. Honest mistakes reported quickly are treated as a chance to improve the policy.

12. Review date

[NAME] reviews this policy and the register **at least every [12] months**, and sooner if we add a major new tool, a tool changes its data terms, the law changes, or something goes wrong. The next review is due on **[DATE]**.

13. Staff acknowledgment

Every person covered by this policy must read it and sign the **Staff Acknowledgment Form** before using AI tools for company work, and again after each major update. Signed forms are kept [LOCATION].

Appendix: settings to check before approving a tool

Guidance for the owner. Use this when you fill in the register. It lists the settings and plan types for four widely used providers, as their help pages showed them on Oct 24, 2026. Providers change their plans and settings often, so check the tool's current privacy page every time you review the register. Mentioning these products doesn't mean we recommend them over others.

Tool	What to check	Where to look
ChatGPT (OpenAI)	On personal Free, Plus and Pro accounts, chats can be used to train OpenAI's models unless "Improve the model for everyone" is turned off. Content from ChatGPT Business, Enterprise and Edu workspaces, and from the API, isn't used for training by default. Temporary chats aren't used for training and may be kept for up to 30 days. Thumbs-up/down feedback can send the whole conversation for training.	Settings > Data controls
Claude (Anthropic)	On personal Free, Pro and Max plans, chats are used to improve models if you allow it in privacy settings, or if they're flagged for safety review. Inputs and outputs from Anthropic's commercial products (such as Claude for Work and the API) aren't used for training by default. Feedback you send is stored for up to 5 years.	Privacy settings; for team plans, Organization settings > Data and Privacy
Gemini app (Google, personal account)	With "Keep Activity" on, Google uses chats to improve its services, including training, and some chats are read by human reviewers; reviewed chats are kept for up to three years. Activity auto-deletes after 18 months by default (can be changed to 3 or 36 months, or set not to auto-delete). With Keep Activity off, chats are kept for 72 hours and not used for training unless you send feedback. Don't enter confidential information.	Gemini Apps Activity, in your Google Account
Gemini in Google Workspace (business account)	Prompts, Workspace content and responses aren't used to train generative AI models without your organization's permission, and content isn't reviewed by humans outside your domain without permission. Interactions stay within your organization. Admins can set conversation retention.	Google Workspace Admin console; Generative AI in Google Workspace Privacy Hub
Microsoft Copilot (work or school account; formerly Microsoft 365 Copilot)	Prompts, responses and data accessed through Microsoft Graph aren't used to train the foundation models, and Copilot shows only data each user already has permission to see. Tidy up file-sharing permissions first , because Copilot can surface anything a user can open.	Microsoft 365 admin center; Microsoft's "Data, Privacy, and Security for Microsoft Copilot" page
Microsoft Copilot app (personal Microsoft account)	Different rules from the work version. In the older app, conversations from signed-in personal accounts can be used to train Microsoft's AI models unless you opt out, and conversation history is kept for 18 months. A new version of the app came out on Aug 18, 2026 with different settings, so check the app's current privacy settings. Approve only the work or school version for company information.	Copilot app Settings (older app: Privacy > Training on conversation activity); Microsoft's "Microsoft Copilot for individuals: your privacy controls and choices" page
[OTHER TOOL]	[Training setting / retention / human review / business plan?]	[Settings location]

Settings checked by _____

Date checked _____