

A SUPERWISE® LEARN KIT · OCTOBER 2026

AI acceptable use policy template

A fill-in policy for a small or mid-size business whose staff already use ChatGPT, Claude, and the AI features inside everyday software. With an employee quick-reference and a rollout checklist.

A template, not legal advice. Have counsel review the completed policy before you adopt it.

START HERE

How to use this template

An AI acceptable use policy tells people which AI tools they may use for work, which information stays out of them, who checks the output, and who to call when something goes wrong. Fill in every bracketed field, delete the blue drafting notes, and have counsel review the result.

- 01 Name one owner.** A person or role who approves tools, answers questions, and keeps the policy current. A blank owner leaves every rule without a next step.
- 02 List the AI already in use.** Ask each team what they use, and check expense reports, browser extensions, and the AI features switched on inside your email, document, and CRM software. The approved list starts from what people do today.
- 03 Fill in the brackets.** Fields look like [\[this\]](#) . The drafting notes point to the rules that most often apply to accounting, law, healthcare, real estate, and insurance firms.
- 04 Have counsel review it.** Your client contracts, industry rules, state law, and employment practices decide what the final wording must say. This template cannot see any of them.
- 05 Roll it out and back it with a control.** Work through the rollout checklist, hand out the quick-reference, and put a technical check behind the data rules so they hold on a busy day.

IN THIS KIT	WHAT IT IS FOR
The policy	Ten sections to adopt and file, with a signed acknowledgment on the quick-reference. Built on the same six core rules as the free template on superwise.ai .
Employee quick-reference	One page for the wall or the onboarding packet.
Rollout checklist	The steps from draft to signed acknowledgments, with an owner and a date for each.
What a control looks like	How to enforce the data rules technically, and how to test that they hold.
Sources	The frameworks and rules this template draws on, with dates.

Allow AI, through an approved path. A ban pushes AI use out of sight, where nobody can see what data leaves. This policy assumes the firm permits AI for work and sets the rules for doing it safely. The reasoning is in *Should a company ban AI at work?* at [superwise.ai/learn/should-a-company-ban-ai-at-work/](#).

AI acceptable use policy

COMPANY	POLICY OWNER	EFFECTIVE DATE
[company name]	[name or role]	[date]
VERSION	APPROVED BY	NEXT REVIEW
[1.0]	[name, title]	[date]

1 Purpose

[Company name] permits staff to use AI tools for work through the approved tools listed in section 3. This policy sets out how to do that while protecting client information, company information, and the quality of the work we deliver.

2 Scope

This policy applies to every employee, contractor, intern, and temporary worker who uses an AI tool for [company name], on a company device or a personal one. It covers any tool that generates, summarizes, transcribes, translates, or analyzes content, including:

- **Chat assistants** such as ChatGPT, Claude, Gemini, and Copilot, reached through a website, desktop app, or phone app.
- **AI features inside other software**, including email, word processing, CRM, practice management, and accounting tools, plus meeting recorders and note-takers.
- **Browser extensions, coding assistants, APIs, and agents** that can read company files or act in company systems.

3 Approved tools

Use only the tools and accounts listed below for work. Sign in with the company - provided account; a personal account, free or paid, is not approved for company or client information. Ask [owner] before using any other AI tool, installing an AI browser extension, or connecting an AI tool to company email, files, or systems.

TOOL AND PLAN	ACCOUNT	APPROVED FOR	DATA ALLOWED (SECTION 4)
[tool, plan]	[company account]	[uses]	[general only / listed client data]
[tool, plan]	[company account]	[uses]	[data classes]
[tool, plan]	[company account]	[uses]	[data classes]

DRAFTING NOTE

Approve a tool only once you can name four things about it: where prompts and files are processed, how long they are kept, whether they are used to train models, and the contract term that commits the vendor to those answers. Record them beside the tool. If you handle protected health information, a cloud service that stores or processes it is a business associate and needs a signed business associate agreement (BAA) first.⁵ SUPERWISE's Learn pages on ChatGPT and Claude found BAA coverage only on specific business and enterprise offerings, with none on consumer accounts.¹¹ Check each vendor's current terms.

4 Data that must never go into AI tools

4.1 Never, with no exceptions

- Passwords, API keys, access tokens, private keys, recovery codes, and other credentials. If one is pasted by mistake, treat it as exposed and follow section 8.

4.2 Never, unless [owner] has approved the exact tool, account, data, and purpose in writing

- **Client personal information:** names with contact details, dates of birth, Social Security and other government ID numbers, and client records.
- **Health information** about clients, patients, or staff, including anything covered by HIPAA.
- **Financial data:** bank, card, and account numbers, tax returns, payroll, and client financial statements.
- **Confidential material:** anything covered by a client agreement or NDA, privileged legal communications, employee records, and unreleased company financials, pricing, or plans.

4.3 Allowed in approved tools

Public information, general questions, and your own drafts once names, numbers, and other identifying details are removed or replaced with placeholders such as [CLIENT] or [ACCOUNT]. When you are unsure which list something belongs on, leave it out and ask [owner].

5 Required controls

[Owner or IT lead] is responsible for these controls, and for keeping evidence that each one is in place:

- **Managed accounts.** Approved tools use company accounts with single sign-on or multi-factor authentication, and access ends on the day a person leaves.
- **Vendor settings.** Training on company data is switched off where the plan allows it, and retention is set to [period].
- **A technical data check.** Prompts and files are checked for the data in section 4 before they leave the device or the company network, using [control].
- **Logs.** AI requests and any blocked or redacted content are logged and kept for [period], so the firm can show what happened after the fact.
- **Approval for connections.** Extensions, plugins, connectors, and agents that read company data or act in company systems are approved and recorded before use.

6 Review of output

A person must check AI output before it is sent to a client, published, filed, used in code, or used to make a decision. The reviewer remains responsible for accuracy, confidentiality, tone, and the final action. In particular:

- **Verify facts, figures, quotations, and citations** against the original source. AI tools can produce confident statements that are false, which NIST calls confabulation.²
- **Check for confidential data** in the output before it goes anywhere, including content the tool pulled from other files.
- **Keep people in decisions about people.** AI output may inform decisions on hiring, pay, discipline, credit, coverage, or care; a named person makes the decision, and [owner] approves any such use first.

7 Disclosure to clients

[Company name] tells clients how it uses AI in their work through [engagement letter, terms of service, or client notice]. Before client information goes into any AI tool, confirm that the client's contract permits it. Where a client has asked us not to use AI, or a contract restricts it, we follow that instruction. Label AI-generated content where a client, court, regulator, or publisher requires it.

DRAFTING NOTE FOR LAW FIRMS

ABA Formal Opinion 512 says a lawyer needs the client's informed consent before entering information relating to the representation into a self-learning AI tool, and that boilerplate consent in an engagement letter is not sufficient.³ State bars have published their own guidance, such as California's.⁴ Check your state's rules with counsel.

8 Incidents

Report any of the following to [reporting contact] through [channel] within [time, such as one business hour] of noticing it: data from section 4 entered into an AI tool, use of an unapproved tool with company data, harmful or inaccurate output that reached a client, or a tool behaving in a way nobody expected. Reports made in good faith are welcome, and early reports limit the damage.

- 1 Stop and preserve.** Stop using the tool for that task. Do not delete the conversation; note the tool, account, time, and what was entered, when it is safe to do so.
- 2 Contain.** Rotate any exposed credential at once. [Owner] asks the vendor about deletion where the plan allows it.
- 3 Assess notice duties.** [Owner] and counsel decide whether clients, regulators, or insurers must be told, and by when.

DRAFTING NOTE

Deadlines depend on who you are. HIPAA requires notice to affected individuals without unreasonable delay and no later than 60 calendar days after a breach of unsecured health information is discovered.⁶ Firms covered by the FTC Safeguards Rule, which names tax preparation firms among its examples, must report unauthorized acquisition of unencrypted information about 500 or more consumers to the FTC within 30 days of discovery.⁷ State breach laws and client contracts can add their own.

9 Training

Everyone in scope completes AI training before receiving access to an approved tool, and again every [12 months]. Training covers this policy, the data rules in section 4 with examples from our own work, how to check output, and how to report an incident. Each person signs the acknowledgment on the quick-reference, and [owner] keeps a record of who completed training and when.

10 Ownership, exceptions, and review

[Owner name or role] approves tools and exceptions, answers questions, and records each decision with its date and reason in [location of the exceptions log]. The owner reviews this policy every [6 or 12 months], after any reported incident, and whenever a tool, contract, or workflow changes in a way that affects the rules. Breaches of this policy are handled under [disciplinary policy or handbook section].

DRAFTING NOTE

Named roles, documented policies, training, and incident processes are the practices the NIST AI Risk Management Framework groups under its Govern function.¹ A small firm can meet the intent with one owner and one shared log.

AI at [company name]: the one-page version

You can use AI for work. Use the approved tools, keep client and company secrets out, and check everything before it leaves the building. The full policy is at [\[link or location\]](#).

Use these tools, signed in with your work account

- [\[approved tool and plan\]](#)
- [\[approved tool and plan\]](#)
- [\[approved tool and plan\]](#)

Anything else, including a new browser extension or a meeting note-taker, needs a yes from [\[owner\]](#) first.

Before you paste

- Replace names and numbers with placeholders such as [\[CLIENT\]](#).
- Ask whether the client's contract allows AI on their work.
- Unsure which list it belongs on? Leave it out and ask.

Something went wrong?

Stop, keep the conversation, and tell [\[reporting contact\]](#) at [\[channel\]](#) within [\[time\]](#). If a password or key was pasted, change it now. Reporting early is the right call.

Never paste or upload

- Passwords, API keys, and access tokens. No exceptions.
- Client names with contact details, birth dates, or ID numbers.
- Health information about anyone.
- Bank, card, and account numbers, tax returns, and payroll.
- Anything under a client agreement, an NDA, or legal privilege.

The only exceptions are written approvals from [\[owner\]](#) naming the tool, data, and purpose.

Before you send

- Check every fact, figure, and citation against the source.
- Read it for confidential details the tool pulled in.
- You own the result. The AI tool does not.

Acknowledgment

I have read the [\[company name\]](#) AI acceptable use policy. I understand which tools I may use, which information stays out of them, and how to report an incident.

Name _____ Signature _____ Date _____

Rollout checklist

From a draft to signed acknowledgments. Give each step an owner and a date, and keep this page with the policy as a record of how it was adopted.

PREPARE

STEP	OWNER	DATE
☐ Name the policy owner and a backup.		
☐ Inventory the AI tools, extensions, and built-in AI features already in use, by team.		
☐ For each tool you mean to approve, record where data goes, retention, training use, and the contract term behind them.		
☐ Check client contracts and engagement letters for terms that restrict AI use or require notice.		
☐ If you handle health information, confirm a signed BAA covers each tool that will touch it.		

ADOPT

STEP	OWNER	DATE
☐ Fill in every bracketed field and delete the drafting notes.		
☐ Have counsel review the completed policy.		
☐ Set up company accounts, sign-on, and vendor settings for each approved tool.		
☐ Put a technical check behind the section 4 data rules, then test it with made-up client data.		
☐ Approve and date the policy; file it where staff can find it.		


LAUNCH AND KEEP CURRENT

STEP	OWNER	DATE
☐ Run training, hand out the quick-reference, and add both to onboarding.		
☐ Collect a signed acknowledgment from everyone in scope.		
☐ Test the reporting route with a made-up incident and record who received it.		
☐ Open the exceptions log and the training record.		
☐ Put the next policy review date on the owner's calendar.		

What a control looks like

A policy is an instruction to people. A control is a check the system runs on every request, including the one sent at six on a Friday by someone who skimmed the training.

The data rules in section 4 are the ones that fail quietly. Nobody sees a client's account number leave in a prompt, so a written rule alone gives the firm no evidence either way. A technical control closes that gap at the point where a request leaves the device or the company network, before it reaches the AI vendor.



What to ask of any control

- **Coverage.** Which tools and which data types does it check? A control only governs the traffic that passes through it, so match its list against your section 3 tools and section 4 data.
- **Location.** Where does the check run, and does the content pass through anyone else's servers on the way?
- **Evidence.** Does it keep a record of what was checked and what was caught, that you can show a client or an auditor?

Then test it. Type made-up client details into each approved tool and look at what the check recorded. Common options include enterprise data loss prevention, browser management that limits which AI sites and extensions are reachable, and an AI gateway that sits between staff and the AI tools.

One option: SUPERWISE Sentinel. Sentinel is an AI gateway that runs on the PC. Before a prompt reaches ChatGPT or Claude, it replaces emails, phone numbers, account numbers, and API keys with placeholders, and it logs every request with what was caught. On Windows 10 and 11 it installs without admin rights and covers Claude Desktop, Claude Code, and, through a preview browser extension, ChatGPT and Claude in Chrome or Edge. SUPERWISE states that protection runs on the PC and nothing is routed through SUPERWISE. It is free for 30 days with no card.¹⁰ superwise.ai/sentinel/

No control catches everything. Sentinel and tools like it check supported data types on the traffic that passes through them. Keep the section 6 review and the section 8 reporting route in place alongside any control. More in *Are your AI controls enforced, or just written down?* at superwise.ai/learn/enforced-or-written-down/.

Sources

Primary sources, with the date each was published or retrieved. Rules change; check the current version before you rely on one.

- 1 National Institute of Standards and Technology (January 26, 2023), *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1. The Govern function: documented policies, roles, training, and incident processes. Supports section 10. nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf
- 2 National Institute of Standards and Technology (July 26, 2024), *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*, NIST AI 600-1. Names confabulation, data privacy, and information security among the risks of generative AI. Supports section 6. nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf
- 3 American Bar Association Standing Committee on Ethics and Professional Responsibility (July 29, 2024), *Formal Opinion 512: Generative Artificial Intelligence Tools*. Competence, confidentiality, informed consent, supervision, and fees. Supports section 7. americanbar.org
- 4 State Bar of California, Committee on Professional Responsibility and Conduct (November 16, 2023), *Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law*. Supports section 7. calbar.ca.gov/sites/default/files/portals/0/documents/ethics/Generative-AI-Practical-Guidance.pdf
- 5 US Department of Health and Human Services, Office for Civil Rights (October 2016), *Guidance on HIPAA and Cloud Computing*. A cloud service that creates, receives, maintains, or transmits ePHI is a business associate and needs a BAA. Supports section 3. hhs.gov/hipaa/for-professionals/special-topics/cloud-computing
- 6 45 CFR 164.404, *Notification to individuals*, HIPAA Breach Notification Rule (retrieved October 7, 2026). Notice within 60 calendar days of discovery. Supports section 8. ecfr.gov/current/title-45/subtitle-A/subchapter-C/part-164/subpart-D/section-164.404
- 7 Federal Trade Commission (December 2024), *FTC Safeguards Rule: What Your Business Needs to Know*, on 16 CFR Part 314; and the notification amendment, Federal Register document 2023-24412 (November 13, 2023), effective May 13, 2024. Supports section 8. ftc.gov/business-guidance/resources/ftc-safeguards-rule-what-your-business-needs-know
- 8 AICPA & CIMA (April 29, 2026), *Small Firm Generative Artificial Intelligence (AI) Policy Template*. A members-only template for CPA firms, listed for readers with access; its contents were not reviewed for this kit. aicpa-cima.com
- 9 SHRM, *Generative AI Usage Policy Template* (date not shown). Members-only, listed for HR readers with access; its contents were not reviewed for this kit. shrm.org
- 10 SUPERWISE (retrieved October 7, 2026), product pages for SUPERWISE Sentinel, Sentinel for Windows, and pricing. Source for every product statement in this kit. superwise.ai/sentinel/ · superwise.ai/sentinel/windows/ · superwise.ai/pricing/
- 11 SUPERWISE Learn (reviewed October 2026): *AI acceptable use policy: a copyable template; Should a company ban AI at work?; Where does your data go when people use AI?; Are your AI controls enforced, or just written down?; Is ChatGPT HIPAA compliant?; Is Claude HIPAA compliant?* superwise.ai/learn/

This template is general information for business owners. It is not legal advice and does not create a lawyer-client relationship. Laws, professional rules, and vendor terms vary by state, industry, and contract, and change over time. Have qualified counsel review your completed policy before adoption.