AI POLICY TEMPLATE – 2026-10-27

Thank you for downloading

Thank you for downloading this AI policy template from House of Silver. We hope it helps you put clear, practical guardrails around AI, so your people can use it with confidence instead of guesswork.

The template gives you the headings and the questions to work through. The answers are yours to write, because a policy only works when it reflects how your organization actually operates. Involve the people who will live with it, keep the language plain, and expect the first version to evolve. Adjust the scope of the policy to fit your organization. A smaller firm might only need a one-page policy.

If you would value an experienced partner along the way, whether to facilitate the conversations, review your draft, or help your teams adopt what you decide, we would be glad to hear from you.

Erik
Founder, House of Silver
hello@houseofsilver.ai | houseofsilver.ai

About House of Silver
House of Silver helps organizations adopt AI in ways that strengthen their people, keep their data safe, and deliver real impact. This template draws on the method Erik Silver published in "Creating an AI company policy" (The Zweig Letter, issue 1604, October 6, 2025, as Erik Stroemberg).

[Organization name] AI Policy

Version	
Policy owner	
Approved by	
Effective date	
Next review	

Contents

1. Purpose 3
2. Scope and definitions 3
3. Approved tools and access 4
4. Data and confidentiality 4
5. Acceptable and prohibited use 5
6. Higher-risk uses 5
7. Quality and human review 5
8. Intellectual property 6
9. Incidents and escalation 6
10. Laws and regulations 7
    10.1 European Union 7
    10.2 United States 7
    10.3 All regions 7
11. Roles, training and support 8
12. Review and change 8

How to use this template: each section has a heading, questions to answer and points to consider, in italics. Some questions include an example sentence to get you started; adapt it or replace it. Write your policy text under each heading, then delete the italic guidance.

1. Purpose

Questions to answer

- *Why are we writing this policy now, and what prompted it?*

  > ***Example:*** *[Organization name] uses AI tools to support our work. This policy sets out how we use them safely and responsibly, so that our people can work with confidence and our clients can keep trusting us with their information.*

- *What do we want AI to do for our people, and for our customers or clients?*
- *Which three to five principles should guide every decision about AI?*

  > ***Example:*** *We use AI to support people, not to replace their judgment. We protect client and personal information at all times. We are open about when and how we use AI.*

- *What tone should the policy set: enabling, cautious, or somewhere in between?*

Things to consider

- *Start with intent before rules. People follow a policy when they understand what it is for.*
- *Rules that are too rigid stifle useful experimentation, and rules that are too vague add risk. Aim for clear expectations that still leave room to try new things safely.*

2. Scope and definitions

Questions to answer

- *Who does this policy apply to: employees, contractors, partners, interns?*

  > ***Example:*** *This policy applies to all employees, contractors and temporary staff who use AI tools in their work for [Organization name], on any device or account.*

- *What counts as AI for this policy: chat assistants, AI features inside existing software, custom tools, agents?*

  > ***Example:*** *In this policy, AI tools means any software that generates text, images, audio, code or analysis in response to instructions, including chat assistants, AI features built into software we already use, and automated agents.*

- *Does the policy cover personal accounts and personal devices used for work?*
- *Which terms need a plain-language definition so everyone reads the policy the same way?*

Things to consider

- *Scope broadly from the start. Chat assistants, image generators, meeting notetakers and AI features built into existing software all belong here, for internal and external work alike.*

3. Approved tools and access

Questions to answer

- *Which AI tools are approved today, and for which roles?*
- *Do people use company accounts only, or are personal accounts ever acceptable?*
- *How does someone request a new tool or have a suggestion of a new tool, who decides, and how quickly?*

  > ***Example:** To request a new AI tool, contact [role or team]. Requests are reviewed within [number] working days against our data, security and cost requirements.*

- *Which AI features are switching on inside software we already use, and who reviews them before they are enabled?*
- *Who approves a new AI agent before it is put to work, and who decides when one is retired? Is that one person for the whole organization, one person per department, or a small group?*

  > ***Example:** [Each department head] approves new AI agents for [their team], and [role] confirms what each agent may access before it starts work.*

- *How will we keep track of every AI agent in use, what it is for, what it can reach and when it is next reviewed, including the ones teams set up themselves?*
- *Which systems may our AI tools connect to, and who approves a new connection?*

Things to consider

- *Keep the list of approved tools outside the policy, in a separate document you can update often. Tools change faster than policies, and the policy should not need reopening every time one does.*
- *Keep the record of AI agents apart from the list of approved tools. The tools list tells people what they may use. The agent record shows where AI agents extend your teams and what work they do.*

4. Data and confidentiality

Questions to answer

- *Which information must never be entered into an AI tool?*

  > ***Example:** Never enter: [personal data about employees or customers, confidential client information, passwords and access details]. Approved tools only: [internal documents, draft deliverables]. Free to use: [publicly available information].*

- *What have we promised clients and customers about how we handle their information, and does AI use change that?*
- *Who checks tool settings such as data retention and whether our data is used for training?*

Things to consider

- *This is usually the most important section. Draw a clear line between public AI tools, which may learn from what you enter, and secure environments you control.*
- *A simple three-level guide (never, approved tools only, free to use) helps people decide in the moment.*

5. Acceptable and prohibited use

Questions to answer

- *Which uses do we actively encourage? Give two or three examples per role.*

  > ***Example:*** *We encourage staff to use approved AI tools to draft first versions of emails and documents, summarize long material, brainstorm ideas, and check writing for clarity.*

- *Which uses need approval or an extra review before they start?*
- *Which uses are not allowed, and why?*

  > ***Example:*** *Staff must not use AI tools to make final decisions about people, to create content that imitates a real person without consent, or to process confidential information in tools that are not approved.*

- *When something is unclear, who should people ask?*

Things to consider

- *Examples teach faster than rules. Show what good use looks like for each role.*

6. Higher-risk uses

Questions to answer

- *Which of four levels does each AI use sit at: AI assists a person, drafts for review, acts within set limits, or acts and reports back? Which approval does each level need, and which uses must never go beyond a given level?*
- *Which AI agents or automations may act on their own, for example send messages, change records or make purchases, and who approves them?*

  > ***Example:*** *An AI agent drafts [client reports] for review by [role]. It may [file and tag incoming documents] on its own. It never [sends anything to a client] without approval from a person.*

- *When may AI tools record or transcribe meetings and calls, and how is consent obtained from everyone present?*
- *Is AI used in decisions about people, such as hiring, performance or promotion? What human review is required?*

  > ***Example:*** *AI may support, but never make, decisions about hiring, performance, promotion or pay. A named manager reviews and owns every such decision.*

- *Do customers or the public interact with our AI directly, for example through a chatbot? How are they told?*

Things to consider

- *Keep people, not AI, responsible for decisions about hiring, performance and pay. AI can inform those decisions; it should not make them.*

7. Quality and human review

Questions to answer

- *Who is accountable for work produced with AI assistance?*

  > ***Example:*** *The person who uses an AI tool is responsible for the result, just as if they had produced it without AI.*

- *Who answers for each AI agent we run: what it may do, and what it produces?*

> ***Example:*** *Every AI agent we run has a named owner who approves what it may access and answers for its output.*

- *What must be checked before AI-assisted work reaches a client or customer, or informs a decision?*

  > ***Example:*** *Before AI-assisted work is shared outside [Organization name] or used in a decision, [the author or a second reviewer] checks it for accuracy, sources, confidentiality and tone.*

- *When should clients, customers or colleagues be told that AI was used?*
- *What happens when an error in AI-assisted work is discovered after the fact?*

Things to consider

- *Treat AI output like the work of a capable junior colleague: often polished, sometimes wrong. Require review at every level of seniority, and explain why, so newer staff build the judgment that good review depends on.*

8. Intellectual property

Questions to answer

- *Who owns work created with AI assistance: the organization, the individual, or the client?*
- *How do we avoid passing off material that belongs to someone else?*
- *What do our client, customer and supplier agreements say about AI use, and do they need updating?*

Things to consider

- *Prompts and agent instructions are work product. Treat them with the same care as your other intellectual property.*

9. Incidents and escalation

Questions to answer

- *What counts as an AI incident, for example a data leak, a harmful output, or wrong information reaching a customer?*
- *Who do people tell, how, and how quickly?*

  > ***Example:*** *Report any suspected AI incident to [role or channel] as soon as possible, and within [24 hours].*

- *How quickly can we pause or stop an AI agent, and who is allowed to?*
- *How will we record what AI did and why, so any result can be traced back?*
- *How do we make reporting blame-free, so issues surface early?*

Things to consider

- *Align the consequences of misuse with your existing IT and HR policies, so AI use is held to the same standard as the rest of the work.*

10. Laws and regulations

The summaries are general information as of October 2026, not legal advice. Keep the subsection that applies to you, keep both if you operate in the EU and the US, and add any other countries you operate in. Confirm current requirements with legal counsel.

10.1 European Union

The EU AI Act is the main AI law in the European Union, amended in 2026 by the Digital Omnibus. It applies to organizations that use AI in the EU, not only to those that build it.

Already in effect: certain AI practices are banned outright (since February 2025), and organizations must take measures to support AI literacy among staff who use AI. Transparency rules for chatbots and AI-generated content apply from August 2026, with a grace period to December 2026 for some existing systems.

Coming next: obligations for high-risk uses, including AI used in recruitment, performance evaluation and other decisions about employees, apply from December 2, 2027. The GDPR applies whenever an AI tool processes personal data.

Questions to answer

- *Do we operate, sell or employ people in the EU?*
- *Do any of our AI uses fall into the high-risk categories, for example hiring or evaluating people?*
- *How will we show that we support AI literacy among the people who use AI?*
- *Which AI tools process personal data, and do our GDPR records and data processing agreements cover them?*

10.2 United States

There is no comprehensive federal AI law by October 2026. Existing federal law still applies: the FTC acts against unfair or deceptive AI practices and unsupported claims about AI, and anti-discrimination laws apply to AI used in employment decisions.

State and city laws differ by location. For example, New York City requires bias audits of automated hiring tools, Illinois regulates AI in employment decisions, California's privacy rules on automated decision-making phase in from January 2026 to January 2027, and Colorado's narrowed AI Act is scheduled to take effect on January 1, 2027.

A December 2025 executive order directs federal agencies to challenge state AI rules, but no federal law has replaced them.

Questions to answer

- *In which states and cities do we have employees, customers or clients?*
- *Is AI used in hiring, promotion or other decisions about people, and do local rules require notices or audits?*
- *Can we back up every claim we make about AI in our marketing and services?*
- *Do state privacy laws give people rights over automated decisions we make about them?*

10.3 All regions

Questions to answer

- *Which other countries or regions do we operate in, and what applies there?*
- *Who tracks changes in AI law and flags when this policy needs an update?*

Things to consider

- *Bring in legal input early. It is far cheaper than rewriting the policy after launch.*

11. Roles, training and support

Questions to answer

- *Who owns this policy and keeps it current?*
- *What training does everyone need before using AI tools at work?*

  > ***Example:*** *Everyone completes [introductory AI training] before using approved AI tools at work, with a short refresher every [12 months].*

- *Which roles need deeper, role-specific training?*
- *Where can people get help day to day, for example champions, a channel, or office hours?*

Things to consider

- *Adoption comes from onboarding, ongoing education and room for discussion. A policy that is only signed is rarely followed.*

12. Review and change

Questions to answer

- *How will the policy be introduced so people actually read it, beyond a single email?*
- *How will people give feedback or suggest changes?*
- *How often will the policy be reviewed, what would trigger an early review, and where is the version history kept?*

  > ***Example:*** *The policy owner reviews this policy every [six months], and sooner if laws, tools or our use of AI change significantly. The current version and its history are kept in [location].*

Things to consider

- *Plan the first review before you publish. A clear baseline that you revise beats a perfect policy that arrives late.*