THE PRIVATE BRIEF · AI & IP

An AI-Use Policy for a Small Business: What It Must Say, and What It Can Leave Out

ANTHONY CLEMENZA · MANAGING PARTNER

September 2026

IN BRIEF

- A workable AI policy for a small company answers four questions: which tools are approved, what information may go into them, what they may and may not be used for, and what a person must check before anything leaves the building.
- It also names one person who owns it and one date to review it, because a policy nobody owns soon stops matching the tools people actually use.
- Most of what large-company AI governance material contains, such as risk grading for every model, formal impact assessments and standing committees, can wait until your company builds or deploys AI at a scale that calls for it.

It is ten to six on a Thursday. Your account manager has a proposal due at nine the next morning, one paragraph that will not come right, and a chatbot open in the next tab on her personal account. She pastes in the client's brief, the pricing notes and last year's results, and the paragraph comes right. Nobody told her not to. Nobody told her anything.

That is how many companies meet AI at work: through a habit, not a decision. The problem is almost never bad intent. It is a capable person on deadline using the fastest thing available, because nobody said otherwise.

You do not need a governance framework to fix it. You need a short document, about two pages, that answers four questions clearly enough that someone on deadline can follow it without asking anyone. Add one named owner and one review date, and you have the whole thing.

The four questions a policy has to answer

1. Which tools are approved. Name them. A free chatbot on a personal account and the business version of the same product are, for your purposes, different products, because they come with different data terms. Some vendors' consumer versions may use your conversations to train their models unless you switch that off in the settings, while the same vendor's business version does not use business data for training by default. List the approved tools, say which kind of account counts, and say what someone should do when they want to try something that is not on the list. The answer should be a quick yes or no from a named person, within a day or two. A flat refusal does not stop people using the tool. It stops them telling you.

2. What information may go in. The cleanest structure is tiers, because a single blanket rule is either too tight to follow or too loose to mean anything. Three is enough:

- **Open:** anything already public, such as published material, marketing copy and public filings. Any approved tool.
- **Internal:** drafts, plans and unpublished financials. Approved business accounts only.

- **Restricted:** customers' personal information, anything covered by a confidentiality agreement, passwords and credentials, health or payment information, and anything a contract says must stay in a particular place. Not into a general-purpose AI tool at all without a named person's approval.

3. What people may and may not use it for. The list of things not to do is shorter and more useful than the list of things you may. A sensible starting list:

- no final decision about a person (hiring, discipline, pay, firing) made on what a tool produced;
- no legal, medical or financial advice produced by a tool for a customer;
- no customer or client information in a personal account;
- nothing published or sent outside the company without a person checking it first.

The first line matters more if you hire in New York City. The city's Local Law 144 bars an employer from using an automated employment decision tool (software that, under the law's definition, substantially helps decide or replaces the decision about who is hired or promoted) unless the tool has had a bias audit within one year of its use, a summary of that audit is public, and the required notices have gone to candidates or employees. If your hiring software might meet that definition, find out before you rely on it.

4. What a person must check. Every factual claim, every number, every quotation and every name, before it leaves the building. Then write down the rule about responsibility, because it is the one that changes behavior: the person who sends it owns it. "The tool wrote it" does not move the responsibility anywhere.

Two lines people leave out

A named owner. One person who keeps the approved list current, answers requests for new tools, and is allowed to say yes. Often that is the operations lead, not the most technical person, because the job is follow-through. Without a name, the policy is a document. With one, it is a process.

A review date. Every six months, and whenever you adopt a new tool. The tools change underneath you: some vendors' terms let them change the terms themselves on notice, and treat your continued use after the change as acceptance. The companion piece on vendor data terms covers what to look for.

What you can leave out

Most of what you will find in large-company AI governance material. At twenty people you do not need a system for grading every AI model by its level of risk, a formal impact assessment for each tool, a standing AI committee, a bias-testing program for models you did not build and cannot inspect, or a classification scheme for AI incidents. Those exist for organizations that build or deploy models at scale. Adopting them at your size produces a document nobody reads, which is worse than a short one people follow.

If you want a reference point, the best-known one is the AI Risk Management Framework from the National Institute of Standards and Technology (NIST), the federal agency that develops technical standards. NIST released it on January 26, 2023, and says it is intended for voluntary use: no business is required to adopt it. It organizes the work under four functions, which it calls govern, map, measure and manage, and in July 2024 NIST added a companion profile for generative AI, the kind of tool that writes text and makes images on request. It is a good structure, worth an hour of your time, and not something a company your size should implement clause by clause. Take the shape and leave the apparatus.

Without a name, the policy is a document. With one, it is a process.

What it looks like on paper

Here is an illustration, a composite rather than a client: a marketing agency of twenty-two people.

1. **Approved tools:** three named products, on business accounts the agency issues. Personal accounts are not approved for any client work.
2. **Information tiers:** the three above, plus one line of its own. Client material covered by a confidentiality agreement is Restricted until someone has checked what that client's contract says about AI, and a one-page log records which clients have agreed.
3. **Not permitted:** client information in personal accounts; an AI-generated likeness or voice of a real person without that person's written permission; anything published to a client's channels without a named person reviewing it.
4. **Checking:** every claim, statistic or quotation in client-facing work is checked against a source, and the reviewer's initials go in the file.
5. **Owner and review:** the head of operations, twice a year, every March and September.
6. **What we tell clients:** one agreed sentence describing how the agency uses AI, so that three account managers do not improvise three different answers.

That is the entire policy. It took an afternoon to write and one all-hands meeting to introduce, and it can actually be enforced, which is the only test that matters.

Introducing it in one meeting

A policy nobody has heard read aloud is a policy nobody follows. Walk through the three tiers using three real examples from your own work. Name the owner, in the room. Then answer the question everyone is actually sitting with, which is whether they are in trouble for the tool they have already been using. If the honest answer is no, say so plainly. That sentence is what turns quiet use into open use.

Then ask what people are already using. You will learn more in ten minutes than an audit would tell you, and the list will be longer than the approved one. The gap between the two lists is where your policy really starts.

A flat refusal does not stop people using the tool. It stops them telling you.

Where a lawyer actually helps

Not in the drafting: you can write these two pages yourself. In two other places. The first is the contracts you have already signed. A customer agreement, a confidentiality agreement or a data processing agreement (the contract that governs how a service provider handles personal information on your behalf) may already limit who can process the other side's information, and where. A policy written without reading them can quietly contradict them. The second is the tools' own data terms, which decide whether your Internal tier is honest: whether "business account" means what your policy assumes it means.

The firm's AI counsel work follows that order: first an inventory of every AI tool your team already uses and what data each one sees, then the policy, then the vendor paper, which means reading the data terms, the training-rights clause and the exit provisions before you sign. The engagements are described on the firm's AI counsel page.

Where to start

This week, send one question to everyone in the company: which AI tools have you used for work in the last month, and on what kind of account? Promise, and mean it, that nobody is in trouble for the answer. The list that comes back is the first page of your policy.

This is general information about writing an internal AI policy, not legal advice about your company's use of AI. Vendor terms and the rules around them change, so check the current versions before you rely on them.

WRITTEN BY

Anthony Clemenza

MANAGING PARTNER

Anthony is the Managing Partner of Clemenza Law Group and serves as general counsel to a wide range of companies and philanthropic organizations. Nearly every essay here began as a real question from a real client: the contract that landed at 4 p.m., the board that meets on Thursday. He writes them the way he would answer you across the table, plainly and with care.

CLEMENZA LAW GROUP

NEW YORK, NY · (212) 365-4875 · CLEMENZALAW.COM

The current version of this essay is kept at clemenzalaw.com/insights/ai-use-policy-small-business