

The One-Page AI Acceptable Use Policy

Fill in the brackets, put a name against it, and circulate. One page people will actually read beats forty pages nobody opens.

PURPOSE AND SCOPE

[Company] encourages the use of approved AI tools to do better work, faster. This page sets out what is allowed, what is restricted, and what is prohibited. It applies to all staff, contractors and vendors acting for [Company], on any device, in any tool.

DATA AND USES

ALLOWED	Public or internal information in approved tools: drafting, summarising, research, analysis, code assistance. A person reviews anything that leaves the company. [add your examples]
RESTRICTED	Confidential company, customer or employee information, and anything under contract or regulation. Allowed only in [named approved tool] , and only with [approval from the risk owner / privacy review] .
PROHIBITED	Secrets, credentials and keys; regulated personal data [name the categories] ; anything you could not explain to the person it belongs to; and any tool not on the approved list, including personal accounts for company work.

APPROVED TOOLS

Current approved tools: **[tool, tier/plan, who administers it]**. To request a new tool, contact **[risk owner]**. Tools are approved on security, data handling, training use, subprocessors and change notice, not on price alone.

HUMAN REVIEW

A named person reviews and approves AI-assisted output before it reaches a customer, a contract, a regulator or a published channel. Internal, low-stakes work needs spot checks only. You remain accountable for anything you send under your name.

ACCOUNTABILITY

AI risk owner: **[name, role]**. They maintain the approved tool list, decide restricted-data requests, and review this page on **[review date]**.

DISCLOSURE

Say when it is AI. Tell customers and colleagues when they are reading, or being assessed by, materially AI-generated output, in the way set out by **[marketing / legal]**.

IF SOMETHING GOES WRONG

Report suspected data exposure, a harmful or badly wrong output, or an unapproved tool in use to **[risk owner / security channel]** the same day. Reporting in good faith is never penalised.

Acknowledgement. I have read this page and agree to work within it. Name: Date:

Adopt it in a week.

Most AI policies die in review. This one is designed to be in force by Friday.

Five steps

1. **Monday — name the owner.**

One executive accountable for AI risk. Not a committee. Put their name in the brackets before anything else.

2. **Tuesday — fill the three rows.**

Allowed, restricted, prohibited, using your existing data classifications. If you do not have classifications, write what you would not want screenshotted.

3. **Wednesday — list the approved tools.**

Including the ones already in use. A tool you do not name is a tool nobody will report.

4. **Thursday — review with legal, security and one sceptical manager.**

One hour, one pass, one page.

5. **Friday — circulate and ask for acknowledgement.**

Set the review date in the calendar before you send it.

What this page deliberately leaves out

It is the first control, not the whole program. The rest of the governance stack covers:

- Vendor due diligence before contracts are signed: training use, subprocessors, model change notice, exit.
- An AI system register, including the shadow AI already running.
- A human-review matrix by output type and consequence.
- An AI risk register with owners and review dates.
- Monitoring after launch, and a tested way to pause a system.

Check where you stand in four minutes. The governance self-check scores the twelve questions behind these controls and names the one to fix first: ericktruong.com/governance/check

Where to go next

- ericktruong.com/governance — the eight controls, and what NIST AI RMF, ISO 42001 and the EU AI Act actually ask of you.
- ericktruong.com/assessment — a wider readiness score across six dimensions, about seven minutes.
- **Guardrails** — my newsletter on AI strategy, governance and transformation, for the people accountable for getting it right.