

Vibe coding policy template for business

<https://prompt-to-prototype.com/governance/acceptable-use-policy/> · as of 2026-10-22

This is a copy-ready AI acceptable use policy template for apps staff build with AI (vibe coding or citizen development), one to two pages long, and it enables building. It says who may build and with which approved tools, which data classes each tool may use, what needs registration and review at each risk tier, the five never-dos, how to request a new tool, and the consequences of breaking the rules. Review it every six months, with HR and counsel.

Who this is for

IT managers, and the HR and legal partners who sign the policy off. You get a fill-in template with notes, the non-negotiable clauses, a 90-minute drafting session, a rollout plan and guidance on consequences. About 10 minutes.

What goes in an AI acceptable use policy?

Copy this into your policy system, or use the DOCX or PDF download on this page, and replace the parts in angle brackets. It's written for builders, whatever their job, not only for developers. Each clause links to the page that holds the detail, so the policy stays short and the rules live in one place. *Prompt to Prototype, Vibe coding policy template for business*, <https://prompt-to-prototype.com/governance/acceptable-use-policy/>, as of October 2026. Licensed CC BY-SA 4.0. Adapting the template for internal use isn't distribution under the licence.

Acceptable use of AI for building apps and automations

Version <n> · Effective <date> · Owner: <enablement lead (a role)> · Next review: <date + 6 months>

Template note (delete before publishing): adapted from Prompt to Prototype, as of October 2026. Not legal advice. Have HR and counsel review this policy, especially clauses 9, 10 and 12, before you publish it; see "Questions to ask your counsel" on the source page.

1. Purpose. We want you to use AI to build tools that make your work better. This policy sets the few rules that keep those tools safe, supportable and owned.
2. Scope. It covers anything you build or change with AI help for company work, or with company data: apps, spreadsheets with AI features, automations, scheduled jobs and agents, on any account. It applies to everyone, including contractors and professional developers building their own tools.
3. Who may build. Anyone who has completed <the builder basics training>. Building an AI feature or an agent also needs <the follow-on training>.
4. Approved tools. Use only the tools on <the approved-tools list>, on your company account. If none fits, request one at <route>; we aim to answer within five working days.
5. Data. Each approved tool is approved for named data classes (Public, Internal, Confidential, Restricted). Use a tool only with data of a class it's approved for. If you're not sure of the class, ask <the data owner or the help route> before you use the data.
6. Register and review. Register your sandbox workspace once (it may be automatic). When you do, you confirm that you'll go to counsel before any sandbox app collects consent-gate data (biometrics, data about anyone under 18, consumers' health or other sensitive data) or handles information about people in Quebec, the EU or the UK. Individual sandbox apps don't need registering until you share one, use data beyond your own work files, make it reachable from the internet, or let it write to anything shared; then register it at <intake form>. Sandbox apps are paused after 90 days unless renewed or registered at a higher tier. An app's risk tier sets the tests, review and support it needs.
7. The five never-dos.
 - (1) Never put restricted data into a tool that isn't approved for it.
 - (2) Never put keys, passwords or tokens in an app's code, front end or prompts.
 - (3) Never build or host a shared app on a personal account.
 - (4) Never give an AI agent write access to live systems without a sandbox run first, a scoped credential and a restore you've

tested.

- (5) Never turn on permission-bypass modes in coding agents.
8. The two floors (every tier, no exceptions, no downgrade).
 - (1) The trifecta rule. No agent or AI feature may combine private data, untrusted content and a way to send data out (the "lethal trifecta"). Remove one leg before it runs, in Sandbox too.
 - (2) No internet reach without company sign-in below Important. Sandbox and Team apps are private by default. An app reachable from the internet without company sign-in must be approved at Important or Critical on an approved public-facing pattern. A builder tool that can't keep a Sandbox app private isn't approved for Sandbox use.
9. Records. Prompts, outputs and agent logs in approved tools are company records under <the retention schedule>. Don't delete them when a legal hold applies.
10. Legal, HR and discipline matters. Don't use a public consumer AI tool for anything about a legal dispute, an investigation or a discipline matter. Use only the approved tool, as counsel directs.
11. Report problems. Report a mistake, exposure or near miss to <single channel> the same working day. Reporting promptly is never a disciplinary matter; hiding a problem can be.
12. Consequences. Breaking this policy is handled first by fixing the app with you. Repeated or deliberate breaches go to your manager and HR under <the disciplinary policy>.
13. Acknowledgement. Read and acknowledge this policy before your first build, and again after each material change.

Notes on the clauses.

- **Clause 3** links to the [training tracks](#). Gate on training, not job title: a finance analyst who did the training is a safer builder than a developer who didn't.
- **Clauses 4 and 5** point to the [approved-tools list](#), which records each tool per data class, and to the data classes on [can I use company data in an AI app?](#). Don't copy the list into the policy; it changes faster than the policy does.
- **Clause 6** follows the [Sandbox rules](#): one workspace registration per builder, with the counsel attestation, and app-level registration from the first answer above Sandbox. It uses the tier model without restating it. Every worker-built app sits in one of four risk tiers (Sandbox, Team, Important or Critical), set by the highest answer to six questions about its data, users, internet reach, dependence, decisions about people and actions, and two floors apply at every tier. The [gate matrix](#) holds what each tier needs; the policy links to it.
- **Clause 9** exists because prompts and outputs can be discoverable, and once litigation is reasonably anticipated you must preserve them; losing them can bring sanctions under FRCP 37(e) in federal court, and states have their own equivalents.¹

- **Clause 10** follows a February 2026 federal district court ruling that a defendant's exchanges with a public AI platform were neither privileged nor work product, partly because the platform's privacy policy allowed disclosure and counsel hadn't directed the use.² It's one court in a criminal case, and it didn't decide anything about enterprise tools; ask counsel how it applies to yours.
- **Clause 11's** no-blame rule is repeated in the [policy on unsanctioned AI use \(shadow AI policy\)](#), which covers unapproved tools, amnesty and enforcement. Keep the two documents separate so this one stays about what's allowed.

Don't: write 'engineers must'

Most of your builders aren't engineers. A policy addressed to developers tells analysts, ops staff and HR that it doesn't apply to them, and those are the builders on AI app builders. Address "you", and define who's in scope.

Don't: put the spend cap in this policy

A per-tool monthly spend cap changes with pricing and contracts. Keep it with the cost controls on [AI spend](#), and let the policy say only that usage-based tools have caps.

What are the non-negotiables?

Clauses 7 and 8 are the only rules that apply at every tier with no exception, no downgrade and no bounded exception. Copy them word for word: the five never-dos have their home on [the builder handbook](#), and the two floors on [the risk tiers page](#). If your version paraphrases them, the next reviewer will find two rules where there should be one.

Everything else in the template is adjustable: the training you require, the request target, the review period. Write those to fit your business, and record why.

Don't: ban and hope (AP12)

A policy that forbids AI tools without offering an approved one doesn't stop use; it hides it. If the approved-tools list is empty when the policy launches, approve one default tool for each common job first. Sources are on [failure patterns](#).

How do you draft it in 90 minutes?

Draft it live, in one session, with the people who'll have to follow it. A policy written by security alone is the one that gets ignored.

Who's in the room: the enablement lead (facilitator), a security lead, a privacy or data-protection lead, an HR partner, counsel or the person who instructs counsel, and at least two builders, one of them a non-developer. Leave vendors out.

Bring: this template, your approved-tools list (even if it has one row), your data classification, and two or three public incidents to ground the discussion.

Time	Topic	Output
0–10	Purpose and ground rules: the policy enables building	Agreed purpose sentence
10–25	Scope and who may build (clauses 2–3)	Scope and the training gate
25–45	Tools and data classes (clauses 4–5)	The first approved tools per class
45–60	Register, review, records, legal matters (clauses 6, 9, 10)	Links filled in
60–75	Reporting and consequences (clauses 11–12), with HR	Draft consequences for HR and counsel to confirm
75–90	Read every clause aloud; test three real cases	Sign-off, or a named blocker and a date

Test it on real cases. "An analyst wants to upload last quarter's customer list to a builder that's approved for Internal data only. Allowed?" If the room can't answer from the text, rewrite the clause.

How do you roll it out?

- **Launch it with something people want:** an approved tool, a training session, a named help route. A policy that arrives alone reads as a crackdown.
- **Tell managers first.** Manager support is a strong lever on how staff use AI,³ and managers get the questions.
- **Track acknowledgement** in the HR system or a simple list, and add the policy to the [builder training](#).
- **Publish a change log** with each version, so people can see what changed and why.

Consequences and review

Not legal advice

Not legal advice. This page summarizes laws as of October 2026 for a general audience. It was written with an AI legal-research persona and hasn't been reviewed by a lawyer for your situation. Ask your counsel.

A written policy with named review dates and stated consequences for non-compliance is the pattern in published public-sector policies; Saskatchewan's May 2026 low-code policy, for example, provides for review, remediation, reassignment to IT or decommissioning of non-compliant apps.⁴ Consequences work when they are:

- **Proportionate:** most breaches are a missing step, fixed with the builder. Clause 12 starts there.
- **Consistent with your disciplinary policy and handbook.** Don't invent a parallel process; point to the existing one, and have HR and counsel read the clause before it's published.

- **Careful about speech rules.** In the US, the National Labor Relations Act makes it unlawful for an employer to maintain work rules that reasonably tend to inhibit employees from exercising their rights to act together about working conditions.⁵ The Board's current test for work rules is *Stericycle, Inc.*, 372 NLRB No. 113 (2 Aug 2023): a rule that has a reasonable tendency to chill those rights is presumptively unlawful unless it advances a legitimate and substantial business interest that a narrower rule couldn't.⁶ It applies to union and non-union employers alike, and covers employees but not supervisors or managers. As of October 2026 the Board may revisit it: in August 2026 the General Counsel asked the Board to overturn *Stericycle* in a pending case, while regional offices keep applying it.⁷ A clause that forbids discussing AI tools, or criticizing them, could be read as chilling. Keep the rules about data and tools.
- **Honest about monitoring.** If the policy says IT monitors AI use, that's employee monitoring, and some states and provinces require written notice or a written policy first. The rules are on the US and Canadian laws page: [state monitoring-notice laws](#) and the [Canadian rules](#), which include Ontario's written electronic-monitoring policy, PIPEDA for federally regulated employers, and the Alberta, BC and Quebec private-sector privacy acts for employee data. In Canada, provincial employment standards also govern how discipline is handled.

Maps to: NIST AI RMF GOVERN 1.1 and 1.2 (legal requirements understood; policies in place); ISO/IEC 42001 clause 5.2 and control A.2.2 (AI policy); SOC 2 CC2.3 and the privacy criteria P1.1 (notice) for the monitoring language. Mapping reuses controls you already have; it isn't certification.

Review dates. Review the policy every six months, and after any incident it didn't cover. Re-check it whenever the approved-tools list adds a new kind of tool, such as the first agent builder.

Don't: publish consequences HR and counsel haven't seen

A consequences clause that conflicts with the disciplinary policy or an employment contract can't be enforced, and the first time you try, everyone learns that it can't. Get both to sign the version you publish.

Depends on your circumstances (C2)

If you're a public body, your policy may be required. US federal agencies must have a generative AI acceptable-use policy that allows use without undue risk (OMB M-25-21).⁸ Check your own jurisdiction's equivalent.

Record the decision

Record the policy's adoption once, using the [decision record template](#).

Illustrative composite

Title: Adopt an enabling acceptable use policy for AI-built apps · **Status:** accepted · **Date:** YYYY-MM-DD · **Owner:** enablement lead (a role) · **Approved by:** HR lead and counsel

Context. Staff build apps and automations with AI tools. Our only AI rule is an email from last year saying "don't paste customer data into chatbots".

Decision. Adopt the 13-clause policy, gated on training, with the five never-dos and two floors copied word for word, and consequences that start with fixing the app.

Alternatives considered. A ban until a full program exists (rejected: bans hide use, AP12). A ten-page policy covering every tool (rejected: nobody reads it, and it goes stale with the tool list).

Consequences. Every builder acknowledges it before building. We report acknowledgement and tool-request times monthly.

Review by: six months from the date above.

Questions to ask your counsel

1. Does our consequences clause match our disciplinary policy and our employment contracts in each state and province where we have staff?
2. Do we need to give notice, or update a written monitoring policy, before we monitor AI use?
3. Could any rule here be read as restricting employees' right to discuss working conditions under the NLRA, under *Stericycle* or whatever test replaces it?
4. How does the February 2026 privilege ruling on public AI tools apply to our approved tools, and should clause 10 say more?

Assumptions and evidence

- **A-004** (Medium): we expect a policy that offers approved tools to bring use into view, where a ban would hide it; test this by comparing discovered unapproved use before launch and three months after.

Evidence: published government policies, survey data on unapproved AI use, US statute and court rules, and one district-court ruling; no controlled study of acceptable-use policies for worker-built apps. Facts on this page are as of October 2026.

Related

- [Shadow AI policy for worker-built apps](#): amnesty, detection and enforcement for unapproved tools.
- [Approve an AI tool or app builder](#): the list and request route clauses 4 and 5 point to.
- [Risk tiers for worker-built apps](#): the tiers and the gate matrix behind clause 6.
- [The builder handbook](#): where the five never-dos live, and what builders do step by step.
- [Can I use company data in an AI app?](#): the four data classes behind clause 5.

-
1. Federal Rules of Civil Procedure, Rule 37(e), failure to preserve electronically stored information. https://www.law.cornell.edu/rules/frcp/rule_37 [statute], accessed 2026-10-21.
 2. *United States v. Heppner*, No. 25 Cr. 503 (JSR) (S.D.N.Y.), memorandum filed 17 Feb 2026 (ECF No. 27). <https://storage.courtlistener.com/recap/gov.uscourts.nysd.652137/gov.uscourts.nysd.652137.27.0.pdf> [court], accessed 2026-10-21.
 3. Gallup, AI indicator: employees with strong manager support are 1.7 times as likely to use AI frequently, May 2026. <https://www.gallup.com/699797/indicator-artificial-intelligence.aspx> [survey], accessed 2026-10-21.
 4. Government of Saskatchewan, Low-Code/No-Code Development and Maintenance Policy, issued May 2026, next review May 2027. <https://taskroom.saskatchewan.ca/-/media/project/taskroom/documents/policy/lowcodenocodedevelopmentandmaintenancepolicy.pdf> [primary], accessed 2026-10-21.
 5. National Labor Relations Board, "Interfering with employee rights (Section 7 & 8(a)(1))", including work rules that reasonably tend to inhibit employees from exercising their rights. <https://www.nlr.gov/about-nlr/rights-we-protect/the-law/interfering-with-employee-rights-section-7-8a1> [regulator], accessed 2026-10-22.
 6. National Labor Relations Board, "Board Adopts New Standard for Assessing the Lawfulness of Work Rules" (*Stericycle, Inc.*, 372 NLRB No. 113), 2 Aug 2023. <https://www.nlr.gov/news-outreach/news-story/board-adopts-new-standard-for-assessing-lawfulness-of-work-rules> [regulator], accessed 2026-10-22.
 7. SHRM, "NLRB General Counsel Targets Stericycle, Other Biden-Era Precedents", 27 Aug 2026 (the General Counsel's position in *Honeywell International Inc.*; the memo doesn't change Board law). <https://www.shrm.org/mena/topics-tools/employment-law-compliance/nlr-general-counsel-targets-stericycle-other-biden-era> [news; secondary], accessed 2026-10-22.
 8. US Office of Management and Budget, Memorandum M-25-21 (generative AI acceptable-use policy), 3 Apr 2025. <https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-21-Accelerating-Federal-Use-of-AI-through-Innovation-Governance-and-Public-Trust.pdf> [primary], accessed 2026-10-21.

Prompt to Prototype, Vibe coding policy template for business, <https://prompt-to-prototype.com/governance/acceptable-use-policy/>, as of 2026-10-22. Licensed [CC BY-SA 4.0](https://creativecommons.org/licenses/by-sa/4.0/).