AI Use Policy Builder

A Simple Starting Point for Small Business

Version: September 2026 | Regulatory references last reviewed: September 2026

Create a practical first-draft AI Use Policy by having a simple conversation with AI.

You do not need to understand privacy law, AI regulation or technical security terms to use this activity. The AI will ask you normal business questions and flag anything that may need a closer look.

What you will do

Tell the AI a little about your business and how you use AI.
Decide what your team can do, what needs caution and what should never be done.
Set simple rules for information, human review and account security.
Receive a first-draft policy, a one-page team guide and a short list of anything that needs attention.

IMPORTANT TO KNOW

This Builder is the first layer of AI governance. Before final rollout, the draft should be followed by a deeper risk review, with qualified professional review where the risks require it.

It is not legal, privacy, HR, cybersecurity, insurance or regulatory advice and it does not certify or guarantee compliance.

You remain responsible for reviewing the final policy, putting the rules into practice and getting qualified advice where your business, industry or AI use requires it.

Laws, standards and AI tools change. Recheck important requirements before relying on them.

BEFORE YOU START

Do not paste real client, employee, financial, health, legal, password or other confidential information into the AI tool.

Use general descriptions instead. Example: “We handle customer contact information and project notes.”

How to use this

You do not need to read or understand the prompt itself. Copy it all, paste it into your AI chat, and let the AI guide you.

1. Copy the full prompt on the following pages and paste it into your AI chat. Then send it. The AI will begin interviewing you one question at a time.
2. Answer in the way that is easiest for you: type in the chat, or tap the microphone icon in the chat box, if your AI tool shows one, and answer out loud. You can switch between typing and speaking at any time.
3. If you are unsure, simply say “I don’t know.” Ask the AI to explain the question or give you a few examples. Do not guess.
4. Before the AI writes your policy, it will summarize what it understood. Correct anything that is wrong or missing.
5. Review the draft and one-page team guide it creates. Before final rollout, complete a deeper AI policy risk review to challenge the draft, identify remaining gaps and determine whether qualified professional review is required.

AI Use Policy Builder Prompt

COPY FROM "PROMPT STARTS HERE" TO "PROMPT ENDS HERE"

PROMPT STARTS HERE

Help me create a practical first-draft AI Use Policy for my small business.

Act as an AI governance, privacy, cybersecurity, risk and business operations advisor, but speak to me like a small business owner, not a lawyer or IT specialist.

Make this easy for me

Ask one question at a time.

Use plain language and short explanations.

Do not give me a long questionnaire.

If I say I do not know, briefly explain why the question matters and give me two or three simple examples or choices based on what you already know about my business. Do not make me guess.

Only ask extra questions when my answer creates a meaningful risk or gap.

Never ask me to paste real confidential, personal, regulated or sensitive information. Ask for categories only.

INSTRUCTIONS FOR THE AI - KEEP THIS IN THE BACKGROUND

Do not ask me to interpret laws, standards or regulations. Do that screening in the background and keep the conversation focused on normal business questions.

Consider relevant privacy, AI, employment, consumer protection, cybersecurity, industry and professional requirements. Where relevant, include PIPEDA and applicable Canadian provincial privacy laws, GDPR, the EU AI Act, California privacy and automated-decision requirements, and ISO/IEC 42001 AI management system practices as a governance reference.

Do not tell me that my business or policy is compliant and do not certify compliance. Do not invent requirements, thresholds, vendor practices, technical settings or legal conclusions.

Grounding rule: do not fill in missing facts or assume what I mean. Aim to have enough information to be at least 95% confident that you understand the relevant facts before making a business, risk or policy recommendation. If you are not, ask the minimum follow-up question needed. If a later answer conflicts with an earlier answer, stop and ask which is correct. Do not resolve contradictions yourself. Do not invent a confidence score.

For current laws, regulations, AI-tool settings, vendor privacy practices, retention, training or security features: verify from reliable current official or primary sources when you have access to them. If you cannot verify an important point, ask me to confirm it or mark it NEEDS EXPERT REVIEW rather than guessing. Do not assign a specific use, information type, permission, account status, integration or control to an individual AI tool unless I have explicitly confirmed it.

If uncertainty remains after reasonable follow-up questions, clearly state what is uncertain, mark it NEEDS EXPERT REVIEW and explain the issue in plain language.

Only surface regulatory detail when it affects a decision I need to make. Connect every regulatory or expert-review flag to a specific answer or confirmed fact. Do not list laws or requirements simply because they might apply.

When to ask a few extra questions

Go a little deeper only if I mention something like:

Personal, employee, health, financial, children's or other sensitive information.

AI helping make important decisions about people, such as hiring, performance, eligibility, pricing, insurance or access to services.

AI giving legal, financial, medical, safety or regulated advice.

AI connected to email, files, CRM, accounting, HR or other company systems.

AI talking directly to customers or creating public-facing content.

Unapproved tools, personal accounts or unclear account ownership.

A regulated industry or customers, employees or operations in more than one jurisdiction.

Ask only what you need to understand the risk. If AI is making or materially influencing important decisions about people, or sensitive information is involved, do not treat the standard first draft as sufficient. Ask the minimum additional questions needed and flag the area for the Advanced Risk Review and/or expert review. If expert review is the right next step, say so instead of turning this into a legal interview.

The guided conversation

Ask these questions one at a time. Use examples when helpful.

1. Tell me about your business

What does the business do? Where does it operate, and where are most of its customers or clients? Ask whether the business is in a regulated profession or industry only if that is not clear.

2. Who is using AI?

Is AI used only by the owner, or by employees, contractors or a wider team?

3. Which AI tools are you using?

Ask me to name the tools I know about. Help me think of common ones I may have missed, such as ChatGPT, Microsoft Copilot, Gemini, Claude, meeting-note tools, design tools or AI built into other software. Do not assume what any individual tool is used for.

4. What are you using AI for?

Ask for the main uses. Examples: writing, marketing, research, sales, customer service, meeting notes, data analysis, HR, hiring, finance, operations, coding, client work, decision support or automation.

5. What kinds of information does your business handle?

Use categories only. Examples: customer contact information, employee information, financial information, health information, contracts, client files, confidential business information, intellectual property or regulated information.

6. What should never go into AI?

Help me sort information into four simple groups: OK TO USE, REMOVE IDENTIFYING DETAILS FIRST, GET APPROVAL FIRST, NEVER PUT INTO AI.

7. Where must a person stay in control?

Ask what decisions or work should always have human review. Help me think about hiring, firing, performance, pricing, contracts, client advice, legal or financial matters, health or safety, compliance and public-facing work.

8. Who is responsible when something goes wrong?

Ask who approves AI use, who the team should contact with questions or incidents, and whether there is a simple way to report a mistake, data exposure or unsafe use.

Do one quick safety check

Before writing the policy, ask only the simple safety questions that are relevant:

What is the current status of each AI tool: APPROVED, RESTRICTED, UNDER EVALUATION or NOT APPROVED?
Are people using business accounts where available, rather than personal accounts?
Is two-step verification (MFA) turned on where available?
Are shared logins being avoided?
Have settings for training, history, retention and sharing been checked?
Does any AI tool connect directly to company email, files, CRM, accounting, HR, customer data or other systems?
Has the business looked at how an AI vendor handles business or personal information?

If an AI tool connects to business systems, ask which tool connects to which system and what kind of access it has, using categories only. Do not assume the mapping. If an important answer is missing, ask the minimum follow-up question needed. Do not turn this into a full risk assessment.

Confirm what you heard before drafting

Before creating the policy, give me a short plain-language summary of what you understand about my business, AI tools and their status, AI uses, information, human-review needs, connected systems, account controls and who is responsible. Then ask: “Is anything incorrect or missing?”

Do not draft the policy until I confirm the summary or correct it. If I identify a gap, ask only the follow-up question needed to resolve it.

Now, create my first-draft AI Use Policy

Keep the policy clear enough that a small business team can actually use it.

At the top of the policy, include a simple policy record: Policy Owner | Approved By | Version | Effective Date | Last Reviewed | Next Review Date. If a field is not known, leave it blank or mark it TO CONFIRM rather than making it up.

1. Purpose

Why we use AI and why responsible use matters.

2. Who This Applies To

Who must follow the policy.

3. Who Is Responsible

Who approves AI use, answers questions and owns the policy.

4. AI Tool Register

Create a simple table: Tool | Status | What we use it for | Who can use it | Connected systems/data | Information rules. Use only these status labels: APPROVED, RESTRICTED, UNDER EVALUATION, NOT APPROVED. Only fill a tool-level field if I explicitly confirmed it; otherwise write TO CONFIRM.

5. What You Can Use AI For

Give practical examples based on my business.

6. Use Caution / Get Approval

List activities that need human review, manager approval or extra safeguards.

7. Never Allowed

List prohibited uses clearly.

8. Information and Privacy Rules

Explain what is OK, what must be anonymized, what needs approval and what must never be entered. Use only the minimum information needed.

9. Human Review

Explain when AI may assist and when a person must check the work or make the final decision.

10. Check AI Output

Explain that AI can be wrong, outdated, biased or missing context and that important work must be verified.

11. Account and Security Rules

Use plain language for approved accounts, passwords, two-step verification, shared logins, file uploads, links, connected systems and suspicious activity.

12. If Something Goes Wrong

Explain who to contact and what to do if confidential information is entered, an account is compromised, an unapproved tool is used or an AI mistake creates risk.

13. Team Expectations

Explain what anyone using AI should know before using it for work.

14. Review and Updates

Make this an ongoing business practice. Set two review points: a simple 6-month check-in and a full review at least every 12 months. Review sooner if the business adds a new AI tool, starts using AI in a new way, connects AI to company systems, handles new types of information, expands into a new jurisdiction, has a significant AI/privacy/security incident, or becomes aware of a relevant regulatory or vendor change. Include the next 6-month check-in date and next annual review date in the policy. Also create two short reminder messages the owner can copy into their calendar.

15. Team Acknowledgement

Create one short acknowledgement employees or contractors can confirm after the policy has been explained to them.

Create a one-page team guide

Keep this short and practical. Use these headings:

OUR AI TOOLS AND STATUS
WHAT YOU CAN USE AI FOR
USE CAUTION / GET APPROVAL
NEVER PUT THIS INTO AI
WHAT AI CANNOT DECIDE
CHECK AI OUTPUT
IF SOMETHING GOES WRONG
WHO TO ASK

Finish with: What needs attention?

Keep this brief. Only report issues supported by my answers. Make each item specific and actionable. If you use a priority, choose: ADDRESS BEFORE ROLLOUT, ADDRESS SOON, or MONITOR / REVIEW. These are action priorities, not legal risk ratings.

READY FOR OWNER REVIEW

Controls and rules that look reasonable as a practical starting point for the owner to review. Do not call the business legally compliant, approved for rollout or fully compliant.

FIX OR CLARIFY

For each practical gap, state: Category | Trigger | Priority | Next action.

GET EXPERT REVIEW

Only list issues that genuinely need expert review. Tie each flag to a specific answer or confirmed fact. For each issue state: Category | Trigger | Why it matters | Priority | Next action | Who should review. Do not invent severity, likelihood, compliance status or legal conclusions. Always mark the affected draft text NEEDS EXPERT REVIEW in yellow highlight when supported, or [NEEDS EXPERT REVIEW] in all caps when highlighting is unavailable.

Do not assume that an issue is legally resolved just because it does not appear on this list.

Offer review reminders only with permission

After the policy and review dates are complete, explain that AI governance is ongoing and recommend a 6-month check-in and a full review at least every 12 months, with earlier review when something important changes.

Then ask me: “Would you like help setting reminders for your 6-month AI check-in and annual AI Use Policy review?”

Do not create, schedule, modify or enable any reminder, calendar event, recurring task or future action unless I explicitly say yes.

If I say yes and the AI tool can create reminders, ask me to confirm the dates and reminder details before creating them.

If I say yes but the AI tool cannot create reminders, give me the recommended dates and short ready-to-copy reminder wording so I can add them to my calendar myself.

If I say no, do not create anything. Leave the recommended review dates in the policy only.

Also ask: “Who should own the AI policy review?” If I have already answered this, confirm the existing policy owner rather than asking again.

PROMPT ENDS HERE

What to do next

Your first draft is a starting point, not a compliance approval. Before your team relies on it:

1. Read it once

Make sure it matches how your business actually works. Correct anything the AI misunderstood.

2. Confirm your AI Tool Register

Confirm each tool's status, account, permitted use, connected systems and information rules.

3. Fix the practical gaps

Deal with anything listed under Fix or Clarify.

4. Review expert flags

If the Builder identifies a legal, privacy, HR, cybersecurity, insurance or regulatory issue, have the right professional review that area.

5. Complete a deeper AI Policy Risk Review

Pressure-test the draft for privacy, data handling, security, tool and vendor risk, human oversight, incidents and regulatory considerations. Where the review identifies an issue requiring legal, privacy, cybersecurity, HR, insurance or regulatory expertise, obtain qualified professional review before finalizing the policy.

6. Finalize and share

Update the policy, explain it to the team and make sure people know the rules and who to ask.

7. Make the review ongoing

Put a 6-month AI check-in and a full annual policy review in your calendar. Review sooner if your AI tools, uses, data, systems, team, vendors, jurisdictions or applicable requirements change.

THE EVALU8-EVOLVE FLOW

1. AI Use Policy Builder - create your practical first draft.
2. AI Policy Risk Review - pressure-test the draft and identify remaining risks or gaps.
3. Professional Review - bring in legal, privacy, cybersecurity, HR, insurance or regulatory expertise where required.
4. Finalize & Team Rollout - update the policy, explain the rules and put them into practice.
5. Ongoing Review - keep the policy current as AI use and requirements change.

Keep your AI policy current

AI tools, business practices, vendors and requirements change. Treat your AI Use Policy as an ongoing business control, not a one-time document.

Assign an owner: One person should be responsible for the check-ins and keeping the policy current.

Every 6 months: quick AI check-in

Ask: New AI tools? New uses? Changes to data, team, vendors or systems? If nothing material changed, record the check-in and continue.

Every 12 months: full policy review

Revisit the Builder and Advanced AI Policy Risk Review, confirm the AI Tool Register and rules, update the policy, and complete any expert review still needed.

Review sooner when something important changes

A new AI tool or vendor is introduced.
AI is used for a new purpose or important decision.
AI is connected to company email, files, customer data, HR, accounting or other systems.
The business handles new personal, sensitive or regulated information.
The business enters a new jurisdiction or regulated market.
An AI, privacy or security incident occurs.
A relevant law, regulation, vendor term or data practice changes.

Optional reminders: When the policy is finalized, the AI should ask whether you want help setting the 6-month and annual review reminders. Nothing is scheduled unless you explicitly approve it.

The goal is not to make you an AI compliance expert. It is to put sensible rules in place, understand the risks and know when to get help.