

Build Your AI Use Policy: A Step-by-Step Guide

This kit turns StormWind's AI Use Policy Template (V2) into a policy your organization can adopt. Work through it in order: settle seven decisions with the right people, fill in the template's 15 sections, then roll the policy out with the one-page staff sheet in the appendix.

This is not legal advice. Have your Legal and IT teams review your policy before you adopt it.

Keep the template open beside this guide. Section numbers and names here match V2. Anything in [square brackets] is yours to fill in. The guide is written for the IT lead or manager who will draft the policy, working with Legal, HR and a business lead.

Before you draft

Settle seven decisions in one meeting before anyone edits the template. Most sections depend on them, and a policy drafted first and decided later tends to contradict itself.

Who should be in the room

Who	Why they are there
IT (and Security, if separate)	Approves tools and accounts, sets up connectors, runs the security review, handles incidents.
Legal	Names the regulations that apply, reviews vendor terms, checks the policy language.
HR	Owens training records, the disciplinary link and how the policy reaches staff.
A business lead	Says what staff actually need AI for, so the policy allows useful work instead of only banning things.

Keep it to four or five people. Name one person to hold the pen.

The decisions to make first

Decision	Settle this	Lead
Approved tools	Which AI tools staff may use for work, named by product and plan, not just brand. Which uses each is approved for.	IT
Account type	Work accounts the organization issues, only. A personal account on the same product is not approved, nor is a free tier.	IT

Decision	Settle this	Lead
Restricted data	Which data never goes into an AI tool. Map it to the data classification you already use rather than inventing a new one. Decide whether any approved tool may ever receive restricted data, and who can approve that in writing.	Legal, IT
Governance owner	One named person or team who approves tools and staff-built agents, owns the policy and answers questions.	Leadership
New tool requests	One channel for requests, a turnaround (the template says ten business days), the criteria, and what staff do while they wait.	IT
Connectors and agents	Whether AI tools may connect to email, files or calendars, which connectors are allowed, and who turns them on.	IT
Training	Whether training comes before access, who provides it, and how often staff refresh it.	HR

Write the answers on one page. You will use them in nearly every section that follows.

Section by section through the template

Work top to bottom. For each section: what to decide, what to fill in, and a short example of a good answer. The examples are starting points, so change them to fit. Sections 5, 8 and 11 are new in V2, and are covered after Section 15.

Header: Effective Date, Version, Owner and Last reviewed

- **Decide:** who owns the policy. The header names them as [Policy Owner].
- **Fill in:** [Date], [Version number], [Policy Owner] and [Last reviewed date], which Section 15 needs.
- **Good answer:** "Effective Date: [Date] • Version: 1.0 • Owner: [AI Governance Lead] • Last reviewed: [Date]"

1. Purpose

- **Decide:** whether the policy only restricts, or also encourages safe use. Staff follow a policy better when it says what they can do.
- **Fill in:** [Organization Name].
- **Good answer:** add one sentence. "[Organization Name] encourages staff to use approved AI tools to work faster and better, within the rules in this policy."

2. Scope

- **Decide:** who is covered (employees, contractors, temporary staff, vendors working for you), which devices, and whether AI features inside other software count.
- **Fill in:** [Organization Name]. Keep the list of who is covered the same as in Section 1.

- **Good answer:** "This policy applies to employees, contractors, temporary staff and third parties who use AI for [Organization Name] work, on any device, including personal phones and home computers. It covers standalone AI tools and AI features built into other software."

3. Approved AI Software

- **Decide:** the approved list, the account rule, the request channel, and who makes the call.
- **Fill in:** [Approved AI Tools List URL or location] and [AI tool request form or queue].
- **Good answer:** a list entry. "[Tool name], [plan]. Work account issued by IT only. Approved for drafting, summarizing and analysis of Internal data. Not approved for Restricted data."

4. Shadow AI

- **Decide:** what counts as an AI feature, and what happens to tools already in use on the day the policy starts.
- **Fill in:** whom staff ask "when in doubt": a name, mailbox or channel.
- **Good answer:** "AI features include browser extensions, meeting recorders and note-takers, mobile apps, and AI options added to software you already use. Ask IT before you turn one on. Tell IT about any AI tool you already use for work by [date]. Reporting it by then carries no penalty."

6. Data Protection

- **Decide:** what "Restricted" means, mapped to your data classification, and whether restricted data can ever go into an approved tool. Which regulations actually apply to you.
- **Fill in:** replace [GDPR / CCPA / HIPAA] with the regulations Legal confirms, and delete the rest. Name your internal data governance policy.
- **Good answer:** "Restricted data never goes into any AI tool, approved or not, unless [AI Governance Lead] has approved that tool for that data in writing. Swap real names and numbers for placeholders such as 'Customer A' or '\$X'."

7. Internally Created AI Tools

- **Decide:** the threshold for review. Trained staff will build many agents and reusable prompts. Review the ones that carry risk.
- **Fill in:** [AI tool request form or queue].
- **Good answer:** "An agent, prompt or workflow that one person uses inside an approved tool, with no connectors, does not need review. Submit it to [AI Governance Lead] before you share it with others, connect it to company systems, or let it act without a person approving each step."

9. Vendor and Third-Party AI

- **Decide:** who the "procuring party" is in practice, and what triggers an assessment. Include existing vendors that add AI features to a product you already use.
- **Fill in:** where the vendor AI assessment lives and who signs it off.
- **Good answer:** "An existing vendor that adds AI features that process our data counts as a new engagement under this section. The assessment records, in writing, whether our data is used to train models, where it is stored, how long it is kept, and how we delete it."

10. AI Ethics and Bias

- **Decide:** which uses are high-impact (hiring, performance, pay, customer eligibility), and who reviews AI output in those cases.
- **Fill in:** [Organization Name].
- **Good answer:** "AI may help draft or summarize in hiring and performance work. A named person makes and records every decision about an individual. AI output is never the only basis for that decision."

12. Incident Reporting

- **Decide:** the channel, the anonymous option, who triages, and how fast.
- **Fill in:** [IT incident email], [anonymous reporting channel], your response time in place of [one business day], and a second channel for urgent cases.
- **Good answer:** "Report to [IT incident email] or [phone or ticket queue]. IT acknowledges every report within one business day. If you pasted something by mistake, stop, leave the chat as it is, and report it. IT will tell you what to do next."

13. Responsibilities

- **Decide:** who the [AI Governance Lead] is. The Lead approves; IT advises and runs the security review.
- **Fill in:** a name or role for each row, including the [Policy Owner].
- **Good answer:** "[AI Governance Lead]: approves or declines tool requests and staff-built tools. IT: runs the security review, keeps the approved list, manages incidents. HR: tracks training and handles violations. Legal: reviews vendor contracts and regulatory questions."

14. Policy Violations

- **Decide:** how the policy links to your disciplinary process, and how you treat an honest mistake reported quickly.
- **Fill in:** the name of your disciplinary policy.
- **Good answer:** "A mistake reported promptly is handled as a training matter. Hiding a mistake, or deliberately breaking this policy, is handled under [HR disciplinary policy]." This keeps Section 12 credible. People report what they will not be punished for reporting.

15. Policy Review

- **Decide:** who reviews and how often. V2 says every six months, because tools and features change often. Decide what triggers an early review.
- **Fill in:** [policy owner email] and [Questions contact].
- **Good answer:** "[Policy Owner] reviews this policy every six months, and whenever we approve a new AI tool or connector, a law that applies to us changes, or an incident shows a gap."

Before you issue the policy, delete the template's "How to use this template" box and search for any leftover square brackets.

Three sections new in V2

V2 adds three sections, with the clause text already in place. For each, settle what is below, then replace its brackets.

5. Training and Enablement

A policy staff have not been taught is a policy they will break by accident.

- **Decide:** whether training comes before access, who provides it, and how often staff refresh it.
- **Fill in:** [annually / every six months] and [HR / Training team].

8. Connected AI: Agents and Connectors

A chat sees only what someone pastes into it. An agent with a connector to email, files or a calendar can reach whatever that connector can reach. Written rules do not change that. The principle behind this section, in one line: "Rules can't fence in a connector. Once it's connected, assume your agent can reach everything the connector can. The real limits are your tool's controls: block the actions you don't want (send, change, delete), turn the connector off in chats that don't need it where your tool lets you, and disconnect it when you're done."

- **Decide:** whether AI tools may connect to email, files or calendars, which connectors are allowed, and who turns them on.
- **Then:** tools differ in which controls they offer. For each connector you approve, IT should record which actions can be blocked and where the off switch is, and tell staff.

11. Human Review of AI Output

Section 10 covers decisions about people. This section covers everything else an agent produces.

- **Decide:** your rule on disclosing AI assistance.
- **Fill in:** [say when staff must state that AI helped, for example in client-facing or formal work].

Rolling it out

Set the effective date two to four weeks after sign-off, so staff are trained before the policy applies. Then work through these steps in order.

1. **Sign off.** Legal and IT review the final draft, and leadership approves it.
2. **Announce it.** Send a short message from a business leader, not only from IT. Lead with what staff can do and which tools are approved, then the rules, then where to get help. Attach the one-pager and link the full policy. If you set an amnesty date in Section 4, say it here.
3. **Train before access.** Run training before the effective date. Cover the data rules, how to check output, and the controls in each approved tool. Record who completed it.
4. **Hand out the one-pager.** Fill in your IT contact, then post the appendix where staff open their AI tool: the intranet, the approved tools list and onboarding packs.

5. **Publish a short FAQ.** Answer the questions staff will ask first. Which tools can I use? Can I use my personal account? I already pasted something, now what? How do I request a new tool? Can I connect it to my email? Whom do I ask? Add new questions as they arrive in the first month.
6. **Collect acknowledgements.** Ask staff to confirm they have read the policy, through your HR system or learning platform.
7. **Review it.** Check in at 90 days: the questions asked, incidents reported and tools requested. Then review on the cycle Section 15 sets, and whenever you approve a new tool or connector.

Help with the training piece

Training is the one section a policy cannot deliver on its own. If you want help with it, StormWind's AI Accelerator bootcamp trains staff to use AI safely and well at work.

It runs as six weekly live sessions. Each person builds a safe AI agent for their own work, in the tool your organization approves, and learns to use it well and keep data safe. Guides online support them between sessions. The one-pager in the appendix comes from the same data rules the bootcamp teaches.

The next cohort starts January 19, 2027. To enroll your team, go to bootcamp.stormwindstudios.com or book a call at bootcamp.stormwindstudios.com/book.

Appendix: What never goes into an AI tool

For everyone at [Organization Name] who uses AI for work. Questions? Ask [IT contact name, email or channel].

Before you start

- **Check with your IT team.** Ask what [Organization Name] allows before you put anything from work into an AI tool.
- **Use an approved tool.** Only the tools listed at [Approved AI Tools List URL or location].
- **Sign in with your work account.** It keeps your chats inside the agreement [Organization Name] signed. A personal account sits outside it.
- **Connected apps.** An email or file connection IT approved can read the names and addresses in what it opens. Everything else on this page still applies.
- **No approved account yet?** Practice with made-up text, or on paper, until IT gives you one.

Never put this into an AI tool

- **Customer or personal details.** No names, email addresses, phone numbers, home addresses or account numbers, for customers or colleagues.
- **Regulated records.** Health, financial, HR and legal records stay in their own systems, unless IT has told you in writing that this tool may see them.
- **Passwords, keys and access codes.** No credentials of any kind, and nothing that grants access to a system.
- **Anything under an NDA or not yet public.** Unreleased plans, contracts, pricing and source code need clearance first.
- **Anything our policy prohibits.** If the AI use policy keeps it out, it stays out.

Do this instead

- **Swap real details for placeholders.** Write "Customer A", "[client name]", "\$X". Put the real details back yourself.
- **Describe the situation.** An AI tool needs the shape of the problem, not the file it arrived in.
- **Trim a file first.** Remove identifying columns and pages, or use a short extract.

Check the answer before you use it

- **Read it as if a colleague had drafted it.** AI tools state wrong things with complete confidence. The check is yours.
- **Say where AI helped.** Follow [Organization Name]'s rule on disclosing AI assistance.

When you are not sure

- **Ask before you paste or upload, not after.** A question to IT takes a minute.
- **When in doubt, leave it out.** If you would not email it to someone outside [Organization Name], keep it out of the chat.
- **If something went in by mistake,** stop and report it to [IT incident email]. Reporting in good faith is always the right call.