

AI Use Policy

A plain-English AI policy starter for Australian small and medium businesses

AI can save time, but it can also get things wrong or put information at risk. These simple rules help us get the upside while keeping people in charge.

Document control	Complete before approval
Organisation	_____
Policy owner	_____
Approved by	_____
Effective date / version	_____
Next review date	_____
Where staff report an AI concern	_____

Before you use this with your team

This is a draft to make your own—not legal advice or a badge that says you are compliant. Fill in the blanks, name the person who looks after AI use, and check the rules against your real tools, work, contracts and customers. If AI will use personal or sensitive information, or could affect someone’s job, money, rights, health or safety, get qualified advice first.

Before you put this into practice

- Choose who owns the policy, who signs off higher-risk uses and who staff can talk to if they are unsure or spot a problem.
- Ask the team which AI tools they already use—including features built into software you already pay for.
- Agree what information can go into each approved tool. The traffic lights in section 7 are a simple place to start.
- Decide who needs to review a use that touches customers, staff, money, rights or safety—and who can approve an exception.
- Check that these rules fit how you handle customer and staff information, security, records, contracts and any industry requirements.
- Give the team a short walk-through: which tools are okay, what not to paste in, how to check answers and how to raise a concern.

Contents

1. What this policy is for • 2. Who and what this covers
3. A few useful terms • 4. Our everyday rules
5. Who looks after what • 6. Which tools are okay to use
7. What information can go into AI • 8. Match the checks to the risk
9. Before you buy or connect a tool • 10. Check the answer before you act
11. Try it safely and keep an eye on it • 12. If something goes wrong
13. Keep people informed and review the rules • Appendix A. AI use-case screen
- Appendix B. Supplier check • Appendix C. AI register fields
- First 30 days • References and source notes

1. What this policy is for

AI can help us get a first draft on the page, make sense of a long document or spot patterns in information. It can also make confident mistakes, expose information or affect people in ways we did not intend. We will use it when it solves a real problem—and we will keep a person responsible for the result.

These are the ground rules for choosing, trying and using AI at work. A quick first draft should not need a committee. A tool that can affect a person, handle sensitive information or take action in an important system needs a closer look. The checks grow with the possible impact. [1][2]

2. Who and what this covers

These rules are for everyone doing work for us—owners, staff, contractors and anyone acting on our behalf. They apply whether AI is a separate app, a feature tucked inside software we already use, or a personal account used for a work task.

“AI” here includes tools that write or create things, make predictions or recommendations, answer questions, help with decisions, or take steps on our behalf. A basic spreadsheet formula is not AI; an AI feature inside our spreadsheet, email or customer software is covered.

This policy sits alongside our privacy, security, record-keeping, staff, safety and complaint procedures. If another rule is stricter, follow that one and ask the person responsible.

3. A few useful terms

Term	Meaning in this policy
AI system	Software that uses information to create text, images or code, make a prediction or recommendation, or help with a decision. It might be a separate tool or a feature inside software we already use.
Generative AI	AI that creates something new—such as text, an image, audio, video or code—from an instruction or other information.
AI use case	The specific job we want AI to do, who will use it, what information it will see and what happens to its answer. The same tool can be low-risk for one job and high-risk for another.
AI agent	AI that can use connected tools to do things—not just suggest them—such as send a message, change a record or start a transaction.
Personal information	Information that identifies a person, or could reasonably identify them. An AI guess or profile about someone can count too. Use the definition that applies to our business.
Sensitive information	A category of personal information that has extra protection under privacy law. If we are not sure, we treat it as sensitive and ask for advice.
Confidential business information	Information we would not want shared publicly: for example, customer or staff details, finances, contracts, plans, passwords or how our systems are secured.
AI system owner	The person who looks after a particular AI use—from the reason for using it to the checks, follow-up and review.
Human review	A person with the right knowledge checks the answer, fixes mistakes and makes—or clearly approves—the final call.
High-impact use	A use that could make a real difference to someone’s work, money, rights, safety or access to an important service—or seriously disrupt our business.

4. Our everyday rules

Start with the problem. Know what we want to improve before we reach for a tool.

A person owns the outcome. AI can help, but it cannot take responsibility for what we do with its answer.

Share less; limit access. Only share what the approved tool needs. Keep private or security-sensitive information out of tools not cleared to handle it.

Give bigger risks a closer check. A rough internal draft needs less checking than AI used in a decision about someone's job, money, rights or safety.

Check it before you trust or share it. AI can sound sure and still be wrong, old, unfair or tricked. Check the important bits.

Be fair. Be upfront. Think about who is affected. Explain AI's role when it matters, and give people a way to ask questions or challenge a result.

Know how to stop and carry on. For important work, know how to pause the tool and keep things moving another way.

Learn as we go. Keep a simple note of important uses and problems. If something is not working, change it or stop.

5. Who looks after what

In a small business, one person might cover more than one role—that is fine. Just be clear about who owns each decision. If AI could have a big effect on someone, get a second person or an outside adviser to check it where you can. [1][2]

Role	Minimum responsibility
Business owner / directors	Approve these rules, decide what level of risk the business is willing to take, and hear about serious problems or high-impact uses.
AI policy owner	Keep this policy and the AI list up to date. Help people work out which checks they need and when to get privacy, security, legal, staff or safety advice.
AI system owner	Own the day-to-day use: say what it is for, fill in the quick check, test it, make sure a person checks important answers, watch how it works and know what to do if it fails.
Team leaders and users	Use the approved tool for the approved job. Check its answers, follow the information rules and speak up quickly if something looks wrong.
Privacy / security / specialist adviser	Give specialist advice when the use calls for it. This could be someone inside the business or an outside adviser.
Supplier or developer	Explain what the product can and cannot do, how it handles our information, what changes and what happens if there is a security problem. Put important commitments in writing.

6. Which tools are okay to use

Before trying a new AI tool—or giving an existing tool a new job—fill in the one-page check in Appendix A and get the sign-off in section 8. Add the approved use to the AI list in Appendix C. Remember to include AI features inside software we already use.

Use a tool only for the job and information we have approved. If it starts using different information, reaches more people, connects to another system or can do more on its own, stop and check it again.

7. What information can go into AI

These traffic lights are a simple house rule, not a legal label. The OAIC's best-practice advice is not to put personal information—especially sensitive information—into public generative-AI tools. A paid plan does not automatically make prompts private, so check the settings and terms first. [3][4]

Information level	Examples	Rule for AI use
GREEN — public	Information anyone can already see: published web pages, public product details or material cleared for release.	Okay to use in an approved tool. Check the answer before sharing it, and do not add private background by accident.
AMBER — internal	Everyday business information that is not public but would cause limited harm if exposed—for example, a routine internal draft.	Use only a business tool we have approved. Check who can see what you enter and whether it is kept or used to improve the service. Share only what the task needs.
RED — restricted	Information that could seriously harm someone or the business if it got out—personal or health details, customer and staff records, passwords, security details, confidential contracts or unreleased plans.	Never put this into a public or unapproved tool. Use it only for a specific approved job, in a service checked for the task, and share as little as possible. Never enter passwords or security keys.

If information points to a person, treat it carefully both going into AI and coming back out. Do not use someone's details for a new purpose just because the tool accepts them. Before any approved use, check why the information is needed, whether we are allowed to use it, who else can see it, how long it is kept and what the result could mean for the person. Get a privacy review or qualified advice when the impact could be significant. [3][4]

When you can, take names and identifying details out first. A person might still be recognisable from a combination of details. If you are not sure, treat it as restricted and ask.

8. Match the checks to the risk

Before a new use goes live, answer the quick check in Appendix A. Look at the job—not just the AI brand: what it sees, who could be affected, what it can do and what happens if it gets things wrong. Our risk bands are a house guide, not legal categories. If more than one band fits, use the higher one. The same tool can be low-risk for one job and high-risk for another. [2]

Pause and get another set of eyes if any answer is “yes”

- Could AI influence a decision about someone's job, pay, credit, insurance, eligibility, health, safety, rights or access to an important service?
- Will it see personal, sensitive, confidential or security-sensitive information?
- Might a customer, staff member or member of the public trust the answer as if a qualified person had written it?
- Can it send a message, change a business record, run code, grant access, buy something or move money?
- Could a mistake affect a vulnerable person, lots of people, an important service or the business's ability to keep going?
- Are we unsure what information it uses, who can see it, how we check its answers or how we stop it if things go wrong?

Our risk level	What it could look like and the basic checks	Who says yes
LOW	A first draft or brainstorm using public information, with a person checking it before use. Use an approved tool and add the tool and task to the AI list.	Team leader, or a pre-approved low-risk list. The policy owner records it.

Our risk level	What it could look like and the basic checks	Who says yes
MANAGED	AI uses internal business information or helps with customer-facing work. Check the provider and data settings, limit access, try realistic examples, keep a person in the loop and review how it goes.	The use owner and policy owner; bring in a privacy or security reviewer if needed.
HIGH IMPACT	AI helps with a decision about a person or handles sensitive information, faces the public, or can act in important systems. Write down the purpose and risks, check the supplier, test it, keep a qualified person responsible, give people a way to ask questions and know how to switch to a fallback.	A senior leader signs off after privacy, security and other needed advice. AI must not make the final high-impact decision on its own.
NOT ALLOWED BY DEFAULT	Using AI to break the law, discriminate, deceive people, bypass security or leak restricted information—or letting an agent make an important external change without a clear approved safeguard.	Stop. We cannot approve anything unlawful. A narrow exception to an internal rule needs a clear reason, an end date and written sign-off from the business owner after suitable advice.

Keep the paperwork in proportion. A low-risk use may need only the one-page check. For something high-impact, write down who might be affected, what could go wrong, what information and suppliers are involved, what testing showed, who makes the final call and how you would respond. Recheck it when the tool, information or job changes. [2][6][7]

9. Before you buy or connect a tool

Before paying for an AI tool or connecting it to business information, do the supplier check in Appendix B. Find out what happens to what you enter, who can see it, where it is kept, whether it is used to train the tool, and what the supplier will do if there is a security problem. Read the terms as well as the sales page. [5]

- Use a business account where you can. Turn on multi-factor authentication—the extra sign-in check—and give each person their own account.
- Give the tool access only to what it needs. Before connecting it to another system or uploading files, check what it will be able to see or change.
- If AI can take action, treat it with extra care. Test it safely, keep its permissions narrow and have a person confirm before it sends something outside the business, moves money or makes a change that is hard to undo.
- Never put passwords, sign-in codes or security keys into a prompt. Keep them in the approved secure place.
- Know how to switch the tool off, get your information back, contact the supplier and keep serving customers if the service is unavailable.
- Check again if the supplier changes the service, the data it uses, the permissions or the contract.

10. Check the answer before you act

AI can sound confident and still be wrong. Treat its answer as a first pass, not proof. Before you send, publish or act on it, check what matters:

- Facts, numbers, dates, links and claims against a trustworthy source.
- Whether it answers the real question and fits the Australian context. Watch for made-up evidence or important details left out.
- Whether it is fair, easy to understand, safe to share and respectful to the people it affects.
- That it does not give away private information, copy something we do not have the right to use or mislead someone about how it was made.

If an AI answer could have a real effect on someone, a person with the right authority must look at the facts, understand what the AI did and make or approve the decision. They must be able to disagree with it. Keep a short, understandable note of the decision and why it was made.

If AI plays an important part in a conversation or decision, be upfront when that is appropriate or required. Make it clear when someone is talking to a chatbot, and give them a way to reach a person—especially if the issue is sensitive or they disagree with the answer. Check whether your privacy notice or customer terms need an update. [1][3]

11. Try it safely and keep an eye on it

Start with a small trial. Before you begin, decide what a good result looks like and try realistic examples—including a few where the tool might get it wrong. Note its limits and when to pause it. The bigger the possible impact, the more careful the testing should be. [2][5]

Once it is in use, ask the people doing the work whether it is still accurate and helpful. Notice how often answers need fixing, whether complaints are coming in and whether the cost or reliance is creeping up. Check high-impact uses more often. Test again after a big change, and pause the tool if the safeguards are no longer working.

12. If something goes wrong

Tell the person named on page 1 straight away if restricted information goes into the wrong tool, AI gives a harmful or unfair answer, an agent does something it should not, a supplier reports a breach or someone challenges an AI-influenced decision. Speaking up early gives us a better chance to put it right.

1. Stop or limit the tool, connection or account if you can do so safely.
2. Tell the use owner and policy owner. Follow the business's usual security, privacy, safety, complaints or continuity process.
3. Keep a secure note of what happened: the tool, time, information involved, what it did, who may be affected and what you did next. Do not forward sensitive prompts or records around.
4. Get the right person to assess possible harm, whether anyone needs to be told, and what to say to customers or staff.
5. Write down how it was resolved. Restart only when it is safe, and change the training or safeguards so it is less likely to happen again.

The business owner or incident lead decides—after checking the facts and getting advice where needed—whether a regulator, affected person, insurer, customer or anyone else must be notified. Do not guess whether a privacy incident needs reporting.

13. Keep people informed and review the rules

Training

Before someone uses an approved tool, show them what it can be used for, what not to enter, how to check the answer and how to report a problem. A short team briefing is often enough to start. Refresh it when the tools or rules change, and keep a simple note of who has had the briefing.

Records

Keep the AI list, approvals, checks, important test results and incident notes in the usual business record system. Follow the business's normal retention and privacy rules. Do not keep personal information in prompts or notes unless it is needed and approved.

Exceptions

If there is a good reason to make an exception to these house rules, write down why, who and what information are involved, the safeguards, who owns it and when the exception ends. Get written approval before starting. An exception cannot make unlawful use okay or override a contract or legal duty. Close it when it is no longer needed.

Review

Put a date in the diary to review this policy at least once a year. Review it sooner if there is a serious incident, a big change to an AI tool, new rules for your industry or signs the safeguards are not working. The business owner signs off important changes. [1][2]

Appendix A. AI use-case screen

Fill this in for each new AI job—not just once for the tool. If you do not know an answer, write “not sure” and ask. Unknown does not automatically mean low-risk.

Field	Record
Proposed use and business benefit	_____ _____ _____
Tool / supplier / business owner	_____ _____
People who may be affected	_____ _____
Information entered or produced	Public / internal / personal / sensitive / confidential / other: _____
Where information goes and who can access it	_____ _____
Does the tool make or influence a material decision?	Yes / No. If yes, explain and identify the human decision-maker: _____
Can it act, send, edit, purchase, run code or connect to business systems?	Yes / No. List permissions and human confirmation controls: _____
Supplier and privacy/security checks completed	Yes / No / Not applicable. Evidence or open questions: _____
Test plan, human checks and fallback	_____ _____
Proposed band and approvals	Low / Managed / High impact / Not allowed by default. Approver(s): _____
Review date and monitoring owner	_____ _____

Quick screen — stop and escalate if any answer is yes

- ☐ Could it influence a person's job, pay, finances, rights, health, safety or access to an important service?
- ☐ Does it handle personal, sensitive, confidential, regulated or security-sensitive information?
- ☐ Can it take action in an external or important business system?
- ☐ Could a plausible error affect many people, a vulnerable person or business continuity?
- ☐ Are data use, vendor access, output checking or incident handling unclear?

If yes, use at least the Managed band and ask the policy owner whether a High Impact review is required. If the proposal involves a high-impact decision or material autonomous action, do not start until the written approval and safeguards are in place.

Appendix B. Supplier check

Use these questions when choosing an AI supplier or when an existing tool changes in an important way. Keep the answers with your use-case notes. Spend more time on higher-risk uses.

- ☐ What is the tool good at—and what does the supplier say it is not good at?

- ☐ What happens to the information we enter? Where is it kept, how long is it kept and who else can see it?
- ☐ Can the supplier read our prompts or use them to improve its tools? Can we turn that off?
- ☐ How does the supplier protect accounts and information, and can we see who has accessed it?
- ☐ How will the supplier tell us if something goes wrong or the service changes?
- ☐ Who owns what we put in and what the AI creates? Can the supplier reuse either?
- ☐ Can we get our information back, delete it and leave if we need to?
- ☐ What checks, help and human support are available if the tool gets something wrong?
- ☐ Does the agreement support the privacy, security and record-keeping promises we have to make?

Decision: ☐ Accept ☐ Accept with safeguards ☐ Further review ☐ Do not use

Reviewer / date: _____

Appendix C. AI register fields

A simple spreadsheet is enough for many small businesses. Keep one line for each tool and job where the risks or safeguards are different, and limit who can edit it.

Register field	What to record
Name and supplier	Tool name, supplier, account type and contact.
Purpose and workflow	What job it does, who uses it and who could be affected.
Owner and approver	Who owns it, who approved it, when and at what risk level.
Data and access	What information it uses, where it goes, how long it is kept, what systems it connects to and what it can do.
Safeguards	How a person checks it, what supplier and safety checks were done, what users are told and how to carry on if it fails.
Monitoring	How it is going, when it was last checked, any problems and what we changed.
Exit / retirement	How to switch it off, remove access, get or delete information and keep the work going another way.

First 30 days: a light rollout

- Week 1 — Make a list of AI tools and features the team already uses. Name an owner and note what each tool does and what information it sees.
- Week 2 — Agree which tools are okay and what information can go into them. Check the business account settings and turn on the extra sign-in check.
- Week 3 — Pick one or two useful, low-risk jobs to try. Fill in the quick check, try some examples and decide who will check the answers.
- Week 4 — Give the team a short briefing, write down approved uses, make sure everyone knows who to tell if something goes wrong and set a review date.

After a month, you should be able to say

where AI is being used, who looks after it, what information can be shared, how important answers are checked and who to contact if something goes wrong.

References and source notes

These sources helped shape the starter. Australian Government guidance is the main local starting point; OAIC guidance informs the privacy sections; ACSC guidance informs the small-business security steps. NIST is a voluntary reference—not an Australian legal requirement.

- [1] National AI Centre, “Create an AI policy” and AI Policy Guide and Template. [Open source](#)
- [2] National AI Centre, “Guidance for AI adoption: foundations.” [Open source](#)
- [3] Office of the Australian Information Commissioner, “Guidance on privacy and the use of commercially available AI products.” [Open source](#)
- [4] Office of the Australian Information Commissioner, “GenAI tools in the workplace: balancing protection of personal information and business efficiency” (4 December 2025). [Open source](#)
- [5] Australian Cyber Security Centre, “Artificial intelligence for small business” (published and reviewed 14 January 2026; developed with NCSC-NZ and COSBOA). [Open source](#)
- [6] U.S. National Institute of Standards and Technology, “AI Risk Management Framework.” Voluntary framework. [Open source](#)
- [7] U.S. National Institute of Standards and Technology, “Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile” (NIST AI 600-1). [Open source](#)