TRAIGUN AI

Traigun Pty Ltd

Can I Put This Into AI?

A practical safe-use checklist for confidential workplaces

Use AI. Do not leak trust.

Lead Magnet | Version 0.1 | July 2026

Who this is for

Owners, directors, consultants, contractors, professional services teams, NFPs and public-sector-facing organisations handling confidential, client-sensitive or regulated information.

This guide is not legal advice. It is a practical governance checklist to help teams make better first decisions about AI use. Adapt it to your laws, contracts, industry duties, client obligations and risk appetite.

The 30-second rule

Core principle

Do not choose AI based on convenience. Choose based on data sensitivity, approval, confidentiality, privacy, contractual duties and human accountability.

Before anyone copies text into ChatGPT, Claude, Gemini, Copilot, Perplexity, a meeting bot, a coding assistant, a design tool, a local model or an AI agent, they should ask one question:

Question

Would this cause harm, breach trust, or create legal/commercial risk if it became public, inaccurate, retained, reused, or disclosed?

If the answer is yes, maybe, or I am not sure - stop and use the traffic-light decision tree.

The safe-use mindset

- AI can assist drafting, research, summarising and workflow design.
- AI should not silently replace judgment, expertise, approval or accountability.
- Public AI should not receive confidential, personal, privileged or regulated information.
- Local AI reduces some external disclosure risk, but it is not automatically compliant or secure.
- AI agents and connected tools are higher-risk because they can access, retrieve, modify, send or delete information.

Use this guide for

- staff awareness
- contractor onboarding
- management discussion
- AI policy scoping
- risk register preparation
- client conversations about safe AI use

Traffic-light decision tree

Use this simple classification before entering information into any AI tool.

Zone	Information type	Public AI?	Action
GREEN	Public, fictional or generic information. No personal, confidential, client, employee, privileged, regulated or system-sensitive data.	Usually okay with care	Use AI, then verify output.
AMBER	Internal, client-related, employee-related, financial, legal, operational, de-identified, draft or non-public information.	Approval required	Ask manager/risk owner. Use approved tools only.
RED	Personal, sensitive, privileged, confidential, case-specific, regulated, protected, government-sensitive, credential, secret or NDA information.	No	Do not use public AI. Use approved secure workflow or no AI.

Examples

- **Green:** “Rewrite this public website paragraph in plain English.”
- **Amber:** “Summarise this internal project update or client email draft.”
- **Red:** “Analyse this employee medical certificate, client file, legal advice, case material, private contract or customer database.”

If unsure

Do not paste it into AI. Classify the information, remove unnecessary details, use an approved tool, or escalate.

The “Can I put this into AI?” checklist

Answer these questions before using AI for business work. One “yes” can move the task from Green to Amber or Red.

- [ ] Does it contain personal information, identity details, employee details, client information or customer records?
- [ ] Does it contain health, financial, legal, complaint, investigation, disciplinary, children’s or other sensitive information?
- [ ] Is it confidential, privileged, subject to NDA, commercially sensitive or not yet public?
- [ ] Does it include passwords, API keys, tokens, source code secrets, security details or system architecture?
- [ ] Does it include third-party copyright material, licensed content, client-owned material or proprietary templates?
- [ ] Could a person, client, supplier, contractor, staff member or regulated entity be affected by the output?
- [ ] Will the AI tool store, reuse, train on, disclose or route the data through a third party?
- [ ] Is the AI connected to email, cloud storage, CRM, HR, finance, case systems, messaging apps or publishing tools?
- [ ] Does the final output need expert review, legal review, management approval or recordkeeping?
- [ ] Would the organisation be embarrassed, exposed or in breach if the input became public?

Practical rule

If the task needs confidentiality, accountability or expert judgment, AI may still help - but only inside an approved workflow with human review.

Cloud AI, local AI and agents

Public cloud AI

Good for public, fictional, generic or low-risk work. Not suitable for confidential, personal, privileged, regulated or client-sensitive material unless the organisation has formally approved the tool and use case.

Enterprise or approved cloud AI

Can be suitable where the organisation has reviewed privacy, security, retention, contractual, access-control, logging, data-residency and model-training settings. Approval should be use-case specific, not just tool specific.

Local AI

Local AI can reduce external disclosure risk, but it is not automatically safe. Risks include unapproved software, unsafe model downloads, poor patching, no audit trail, copied files, weak access controls and uncontrolled outputs.

AI agents and connected tools

Agents are higher-risk because they may search, retrieve, summarise, email, upload, edit, delete, publish or trigger workflows. They need strict permissions, logging, human approval gates and testing before use.

AI option	Use when	Avoid when
Public AI	Public/generic tasks	Confidential, personal or client-specific tasks
Enterprise AI	Tool/use case is formally approved	Settings, contract or data handling are unclear
Local AI	Managed, approved, secured and documented	Personal device or unapproved model/software
AI agent	Narrow, logged, permissioned, human-approved workflow	Open-ended access to email, files or systems

Human review before use

AI output should be treated as a draft, not a decision. A competent human must review output before it is sent, published, submitted, relied on, or used for client, staff, management or regulated work.

Reviewer checks

- [ ] Is the output factually accurate?
- [ ] Are claims supported by reliable sources or approved internal material?
- [ ] Has the AI invented facts, law, policy, numbers, citations or obligations?
- [ ] Could the output mislead a client, staff member, regulator, director, supplier or member of the public?
- [ ] Does the output expose personal, confidential, privileged or regulated information?
- [ ] Does it need legal, financial, technical, cyber, privacy or management review?
- [ ] Does this AI-assisted work need to be recorded for accountability, audit or continuity?

Simple approval record

AI tool used	
Purpose of use	
Information category	Public / Internal / Confidential / Personal / Other
Reviewer	
Date	
Approved for use?	Yes / No / Conditions

Accountability

The person or organisation using the AI output remains responsible for the final work. AI assistance does not remove professional, contractual, privacy, security or governance obligations.

Minimum AI safe-use controls

A small organisation does not need a Big 4-style AI program to start. It needs clear rules, ownership and repeatable controls.

1. Appoint an AI risk owner.
2. List the AI tools currently used by staff and contractors.
3. Classify use cases as Green, Amber or Red.
4. Approve a short AI Safe Use Policy.
5. Train staff on what must never go into public AI.
6. Add a contractor AI-use clause to engagements.
7. Maintain a basic AI risk register.
8. Require human review before AI output is relied on.
9. Document material AI use where needed.
10. Review the policy every 6 months or when tools/laws/operations change.

Paid implementation pathway

The full Traigun AI Safe Use Starter Pack can convert this checklist into organisation-specific policy, decision tree, risk register, contractor clause, staff guidance and workflow controls.

Suggested internal ownership

Role	Responsibility
Director / owner	Approve policy, risk appetite and accountability.
Manager	Enforce staff use rules and approval process.
Privacy / risk lead	Review personal/confidential information risks.
ICT / cyber lead	Approve tools, access, logging and security controls.
Contract owner	Ensure contractors follow AI-use clause.

References and next step

This checklist has been prepared as a practical governance resource and should be adapted to each organisation's legal, regulatory, contractual and operational context.

Official guidance reviewed

- Digital Transformation Agency - Staff guidance on public generative AI
- Digital Transformation Agency - Agency guidance on public generative AI
- Office of the Australian Information Commissioner - Privacy and commercially available AI products
- National Archives of Australia - Information management for records created using AI technologies

Call to action

Want this customised for your organisation?

Traigun AI can help convert this checklist into a working AI Safe Use Starter Pack: policy, traffic-light decision tree, risk register, contractor AI-use clause, human review protocol and practical staff guidance.

Recommended first engagement: AI Safe Use Starter Pack - a fixed-scope review and template implementation for confidential or regulated workplaces.

Contact: traigun.org

Disclaimer

This document is general information only. It is not legal, privacy, cyber security, procurement, employment, financial, public-sector or professional advice. Organisations should obtain appropriate advice and approvals before relying on AI for sensitive or regulated work.