

TEMPLATE · REPLACE EVERY [BRACKET] · AIM FOR TWO PAGES, NEVER MORE

AI usage policy — [Company name]

OWNER	[name — one person, not a committee]	VERSION	[1.0]
APPLIES TO	Everyone who works here, including contractors	REVIEWED	[date] · next in 6 months

1 · Why this policy exists

AI tools make us faster and are approved for use here. They also move data outside the firm and produce confident text that is sometimes wrong. This policy exists so everyone can use AI *without* guessing where the lines are. It is deliberately short. If a situation is not covered, the rule is: ask [owner] first.

2 · The approved tools

Work may be done with these tools, on **company accounts only**:

TOOL	ACCOUNT TYPE	APPROVED FOR	NOTES
[e.g. Claude — Team plan]	Business	GREEN + AMBER data	Training on our data: off
[e.g. Microsoft 365 Copilot]	Business	GREEN + AMBER data	Within our tenancy
[tool]	[tier]	[scope]	[notes]

Anything not on this list is not approved for work data. To get a tool added, send [owner] the tool name and what you want to use it for; answer within one week.

3 · The data rules — three bins

BIN	EXAMPLES	RULE
GREEN	Published material, blank templates, general research, code with no secrets	Any approved tool, freely
AMBER	Internal drafts, proposals, process documents	Approved business-tier tools only
RED	Customer/staff personal data · financials · contracts · credentials · NDA material · [anything specific to your firm]	Not into AI tools. Exceptions only with written approval from [owner] for a named tool with a contract that covers it

4 · The rules

R.01 Company accounts only.

Work data does not go into personal AI accounts. Historic use: delete work material — [owner] can help.

R.02 You own your output.

Everything AI-drafted is reviewed by a human before it reaches a customer, a supplier, or a decision. "The AI wrote it" is not a defence.

R.03 Meetings: people first.

AI notetakers join meetings only when everyone present has been told. External calls: ask before recording.

R.04 No new connections without approval.

Plugins, integrations, browser extensions and agents that link AI to our email, files or systems need a yes from [owner] before they are connected.

R.05 Agents get the least access that works.

Any AI that can *act* — send, change, delete, pay — runs on its own restricted account, with a named human owner, and its actions are logged.

R.06 Money has one procedure.

New payees, changed bank details and urgent payment requests are verified by phone to a number already on file — regardless of how convincing the voice, video or email was.

R.07 Accuracy is checked, not assumed.

AI output that includes facts, figures, law or code is verified against a source before use. AI is a drafting tool, not a source of record.

R.08 Tell someone, fast.

Pasted the wrong thing? Suspicious AI-flavoured email? Agent did something odd? Tell [owner] the same day. Reporting is never punished.

5 · What happens if this goes wrong

Honest mistakes reported quickly are process failures — we fix the process. Deliberately feeding RED-bin data to unapproved tools, or hiding an incident, goes through the normal disciplinary procedure. The difference between the two is the reporting.

6 · Review

Reviewed every six months by [owner], or immediately when a major tool changes its terms, or after any incident. Suggestions welcome — the approved list is meant to grow.

SIGN-OFF

APPROVED BY	[name, role]	DATE	[date]
COMMUNICATED	[all-hands date · email date · where the cheat sheet is pinned]		

Template © Peter Bassill · peterbassill.com · CC BY 4.0. Adapt freely; keep it to two pages — the length is a feature. This template is general guidance, not legal advice.