

MADE-UP EXAMPLE

Example Studio is not a real organisation. Every name, role, tool and date in it is invented, to show what a finished policy looks like.

Example Studio is a made-up design and marketing studio of about 30 people, with clients in the UK and the EU. Read it next to the template: every bracket has been filled in or deleted.

Example Studio AI use policy

Version	1.0
Approved by	Managing Director
Approved on	3 November 2026
Policy owner	Operations Lead, ops@example.com
Next review	3 May 2027

AT A GLANCE

- **Use approved tools only**, signed in with your work account.
- **Keep these out of every AI tool:** personal data, client confidential material, passwords and security details, unreleased financials and board papers.
- **A named person checks** anything AI helped to make before it leaves the team.
- **AI agents ask first** before they send, pay or delete, and that is set in the tool.
- **Report a slip the same day** to the Operations Lead or the #ai-help channel.

1 Purpose

This policy helps everyone at Example Studio use AI tools to do better work, while protecting our clients, our people and our information. It says which tools we use, what never goes into them, who checks what they produce, how far AI agents may act on their own, and how we stay open about it.

AI tools change quickly, so this policy will too. If something here gets in the way of good, safe work, tell the policy owner (see "Owner, review and questions").

2 Who and what it covers

People. Everyone who works for or with Example Studio: employees, contractors, temporary staff, interns and anyone else who uses our accounts, devices or information.

Tools. Any tool that uses AI to write, summarise, translate, transcribe, analyse or create text, images, audio, video, code or data. That includes:

- chat assistants, and assistants that take a goal and do a whole task, such as researching, writing documents or working across our apps;
- AI features inside software we already use, such as email, documents, spreadsheets, meeting and design tools;
- meeting recorders and note-takers;
- browser extensions, plug-ins and connectors that link an AI tool to other apps;

MADE-UP EXAMPLE / EXAMPLE STUDIO IS NOT A REAL ORGANISATION

- AI agents that can take actions, such as sending messages, booking, buying, browsing or changing files, including agents that run on a schedule.

Places. Work devices and personal devices alike, whenever work information is involved.

Not covered. Nothing at present.

3 Approved tools

The register

We use only the AI tools listed in our AI tools and connections register for work. For each tool it records the plan we use, who may use it, the information allowed in it, the apps it is connected to and whether it can change them, its approval setting, the level its agent tasks may reach, its memory, any scheduled tasks, its owner and its next review date.

The Operations Lead keeps the register up to date and tells everyone when it changes. Where it is kept: the Policies folder on the shared drive.

Accounts

Always sign in with your work account. Never use a personal or free account for work: we cannot see or control how a personal account stores or uses what you put into it.

Asking for a new tool

Anyone can ask for a new tool, a new AI feature or a new connection. This is how it gets approved:

- 1 **Ask.** Fill in the new-tool request form and send it to the #ai-help channel: the tool, the job you want it for, and the information it would touch.
- 2 **Check.** The Operations Lead asks the maker the questions on our vendor checklist: where our information would be stored, whether it would be used to train the tool, whether we get a work account we control, whether we can delete our information, how it connects to other apps, its security and its cost.
- 3 **Try.** If it passes, a small trial with everyday information only (nothing from "What never goes into an AI tool"), for two weeks.
- 4 **Decide.** The Managing Director adds it to the register with what it is approved for, or explains why not.

You will get an answer within 10 working days. Until then, please do not use the tool for work. New AI features that appear in software we already use count as new tools: they stay switched off until the Operations Lead has checked them.

4 What never goes into an AI tool

Never put any of these into any AI tool, whether you type it, paste it, upload it, record it or connect the tool to it:

- **Personal data:** anything that identifies a person, such as names with contact details, ID numbers, health information, HR records and client contact lists.
- **Client confidential material:** anything a client has shared with us in confidence, and anything covered by a confidentiality agreement or contract.
- **Passwords and security details:** passwords, access codes, keys and anything else that would let someone into our systems.
- **Unreleased financials:** results, forecasts, budgets, prices and deals that are not public yet.
- **Board papers:** board minutes, papers and plans before they are shared with the whole studio.

The one exception: approved work tools. An item on this list may go into a tool the register approves for that kind of information, on a work plan whose contract says how our information is stored, used and deleted, and only for the purpose the register lists. Never passwords, and never against a client's contract.

If you are not sure, leave it out, take out the details that make it sensitive (names, numbers, places), or ask the policy owner first. Taking names out is not always enough: if the details together could still point to a person or a client, treat it as confidential.

5 The checking rule

Anything AI helped to write or make is checked by a named person before it leaves the team.

"Leaves the team" means it goes to a client, a customer, a partner or the public; it is published; it goes into a contract, a report or an official filing; or it is used in a decision about a person.

Who checks. The person who used the tool is responsible for what it produced. For client work, anything published, and anything legal or financial a second person checks too: the project lead, or the Managing Director for anything legal or financial. "The AI did it" is never the answer to "who checked this?"

What the checker does:

- reads all of it, not just the start;
- checks every fact, figure, name, date, quote and link against a source outside the AI tool;
- makes sure it says what we mean, in our own voice, and that nothing from "What never goes into an AI tool" slipped in;
- asks: would I put my name to this?

When an agent did the work, check what it did, not just what it wrote. Run the five checks: The job (did it do what the brief asked?), The facts (is every name, number, date and quote real?), The steps (how did it get there?), The changes (what is different now?) and The sign-off (would you put your name to it?).

Why. An AI tool writes a likely answer, and likely is not the same as true. Quotes and links can be invented too. Asking the tool "are you sure?" is not checking: check outside the chat. When an AI tells you its work passed a check, ask what was checked.

Keeping a record. For client work, the checker adds their initials and the date to the project tracker.

6 When we say AI made it

Two of these rules match situations named in the transparency rules of the EU AI Act (Article 50 of Regulation (EU) 2024/1689), which apply from 2 August 2026. The rest are our own choices.

- **Images, audio and video that could pass as real.** If we use AI to make or change an image, audio or video that resembles real people, objects, places, entities or events and could falsely appear authentic, we say clearly that it was made or changed with AI. For clearly artistic, creative, satirical or fictional work, we say so in a way that does not spoil the work.
- **Published text on matters of public interest.** If we publish text that AI generated or changed, to inform the public on matters of public interest, we say so, unless a person has reviewed or edited it and the Creative Director holds editorial responsibility for it. Our choice: such text always gets that review, and the Creative Director signs it off.
- **With clients.** We tell clients when AI tools did a large part of work we deliver to them.
- **AI that talks to people.** We do not use an AI assistant that talks to clients or the public. If we start, it will say it is an AI at the start.

- **Inside the team.** Say when AI did a large part of a piece of work, for example in a line at the end or in the handover note.
- **Real people.** Never use AI to make something pass as a real person's words, voice or face, and never present an AI image, audio or video of a real person as real.
- **When asked.** If someone asks whether AI was used, give a straight answer.

7 AI note-takers and recordings

- **Tell everyone first.** Before an AI tool records, transcribes or summarises a meeting or a call, tell everyone in it, in the invitation and again at the start.
- **Follow our consent rule:** everyone agrees first, and recording stops if anyone objects.
- **Some meetings are never recorded:** HR, disciplinary, health and legal conversations.
- **What we keep, and for how long:** notes are kept in the project folder for 12 months; recordings and full transcripts are deleted after 30 days.
- **Check the notes before they are shared.** The person who shares them checks them first, under "The checking rule".

8 Agents and connections

An AI agent is an AI tool that acts, not just writes: it can send, book, buy, edit files, browse or run on a schedule. These rules apply to every agent and every connection.

- **Every agent task gets a level,** recorded in the register: 1 Read, 2 Draft, 3 Ask first, 4 Alone. New tasks start at level 1 or 2. A task moves up one level after five clean runs, and down at once after a surprise.
- **Three questions set the ceiling:** Can it be undone easily? Does it leave the team? Does it read outside content? If it cannot be undone easily, or it leaves the team, it goes no higher than level 3, Ask first.
- **"Ask first" is set in the tool, not typed in the brief.** Where the tool has an approval setting, we set it so the agent asks before it sends, pays, deletes or changes a shared record. A sentence in a brief is not a lock.
- **Outside words, no power.** An agent that reads outside content (the web, outside email, files others sent us) may not also send, pay or delete without asking.
- **Scheduled agents read, report or draft.** A task that runs on a schedule or starts by itself goes no higher than level 2, Draft. It never sends, pays, deletes or changes existing files.
- **A named owner.** Every agent and every connection has an owner in the register, who reviews its runs weekly for the first month, then monthly.
- **The power to stop it.** The owner knows how to pause, disconnect or switch off each agent. Anyone who sees an agent do something odd stops it if they can, and tells the Operations Lead or the #ai-help channel the same day.
- **Connections.** Do not connect an AI tool to email, files, calendars, chat or other systems unless the policy owner has approved that connection and recorded in the register whether it can only read or can also write.

9 AI in decisions about people

- **A person decides.** AI never makes a decision about a person on its own: hiring, pay, performance, promotion, discipline, dismissal, the work people are given. A person makes the decision and can explain it.
- **Ask before using AI on people decisions.** Before anyone uses an AI tool to screen applications, rank candidates, allocate work, or monitor or score performance, the Operations Lead checks with our adviser first and records the decision in the register.
- **Tell people.** We tell people when AI is used in a decision about them, and how to ask for a person to look again.

- **No emotion reading.** We never use AI to read or guess people's emotions at work.

10 Security and accounts

- **Work accounts only,** with two-step sign-in switched on.
- **Access.** The Operations Lead decides who can use each tool, and removes access on someone's last day.
- **Settings.** For each approved tool, the policy owner checks its data settings, including whether what we put in is used to train it, its memory and its approval setting, and records the choices in the register.
- **Extensions and plug-ins.** Do not install AI browser extensions, plug-ins or apps that are not in the register.

Outside text is untrusted. Words planted in something an AI tool reads, such as a web page, an email or a document, can get treated as an instruction. This is called prompt injection. So we:

- treat what an AI tool reads from outside as untrusted;
- give each AI agent only the access its job needs;
- make it ask before it sends, pays or deletes, set as an approval step in the tool, not a typed request (see "Agents and connections").

These steps cut the risk. They do not end it.

Report a slip the same day. If something from "What never goes into an AI tool" went into an AI tool, or a tool or agent did something you did not expect, tell the Operations Lead or the #ai-help channel the same day. We will not blame anyone for reporting a mistake quickly.

11 Training and the training record

- Everyone completes our AI basics session (about 60 minutes) before using AI tools for work, and a short refresher every six months.
- New joiners complete it within their first two weeks.
- Training covers this policy, the approved tools, what never goes in, how to check what AI produces, how agents and their levels work, and how to report a slip.
- The Operations Lead keeps a record of AI training and support: the date, who took part, what it covered, the materials, how long it took and who led it. Where it is kept: the Policies folder on the shared drive.

Article 4 of the EU AI Act, as amended in 2026, requires providers and deployers of AI systems to take measures to support the development of AI literacy of their staff and of others who operate or use AI systems on their behalf. It does not require them to guarantee any specific level of AI literacy of any individual. This training, and its record, are part of how we do that.

12 If the policy is broken

If something goes wrong, we first look at what happened and fix the cause, whether that is a gap in the tools, the training or this policy. Deliberate or repeated breaches are handled under our disciplinary policy.

13 Owner, review and questions

- **Owner.** The Operations Lead owns this policy and the register.
- **Review.** We review it every six months, and sooner when we add a major tool or agent, when the rules that apply to us change, or after a serious slip.
- **Next review:** 3 May 2027.
- **Questions:** the Operations Lead, or the #ai-help channel.

MADE-UP EXAMPLE / EXAMPLE STUDIO IS NOT A REAL ORGANISATION

- **Sign-off.** Everyone signs the staff sign-off sheet when they receive this policy, and again for each new version.
- **Related policies:** information security, data protection, acceptable use.

Approval

Approved by	Role	Signature	Date
Managing Director	Managing Director		3 November 2026

Version history

Version	Date	What changed	By
1.0	3 November 2026	First version	Operations Lead