

GUIDE AND POLICY TEMPLATE

Guard rails for everyday AI use

What to put in, what to keep out, how to check the output, and how to adapt these rules into your own business policy.

12 sections · about 20 minutes · guide and policy template

For staff using AI at work and the people setting policy. Adapt the examples to your approved tools, information classifications and reporting process.

This reading edition follows the reviewed website content. Examples and exercises remain material for you to adapt to your approved environment.

Contents

- 01 Start here
- 02 Approved and non-approved tools
- 03 What goes in the box
- 04 What to ask, and what to take elsewhere
- 05 Reading the output
- 06 Bias, in the places it bites
- 07 Dos and don'ts
- 08 Who owns the output
- 09 When it goes wrong
- 10 What this is built on
- 11 Adapting this into your own policy
- 12 The one-page card

Start here

This guide is for the person typing the prompt. It assumes no technical background, no legal training and no appetite for a forty-page policy. If you read only the card on the last page you will still avoid most of the ways this goes wrong.

Two things make AI at work different from AI at home. The first is that the material you type in belongs to somebody else: your employer, your customers, your colleagues. The second is that the output goes out under your name, into decisions that affect real people. Everything in this guide follows from those two facts.

Nothing here is a reason not to use these tools. Most of the risk sits in a narrow band of behaviours, and once you know where that band is you can work quickly and confidently everywhere else.

The Shape Of It

Use tools approved for the particular work and information involved. For general knowledge without organisational material, follow your organisation's acceptable-use rules even when using a public tool.

Assume every output is a draft written by a confident stranger who has not checked their facts.

Your name goes on the result, so the last read is always yours.

Approved and non-approved tools

Approval is not a formality and it is not about which product is cleverest. It is about where your words go once you press enter, and who can get them back.

What approval actually buys

Approval should document the following checks for the specific service, tier and use case. An approval label alone does not guarantee any of these protections.

- Whether prompts, attachments and outputs can be used for model training, including any settings or contractual exclusions.
- Where information is processed and stored, how long it is retained, who processes it, and what the contract permits.
- Which permissions the service and its connected tools use, including any shared connections or service identities.
- What records are available, who can access them, and how the organisation meets its retention and disclosure duties.

An unapproved tool may offer protections, but your organisation has not established that they cover this work. Do not assume that a familiar brand, paid account or privacy setting is sufficient approval.

The practical difference

WHAT YOU WANT TO DO	APPROVED TOOL	ANYTHING ELSE
Paste a customer email thread	If approved for this data and purpose	No
Work on an internal draft or report	If approved for this data and purpose	No
Ask about staff, pay, health or performance	Yes, if the tool is approved for it	No
Upload a contract, drawing or data extract	If approved for this data and purpose	No
Explain a concept with no company detail in it	If approved for this data and purpose	If organisational policy permits
Draft something generic you will rewrite anyway	If approved for this data and purpose	If organisational policy permits

The Honest Version

People do not use unapproved tools out of defiance. They use them because the approved one is slower, or was never explained, or nobody told them it existed.

If you find yourself reaching for something else, that is worth saying out loud to whoever runs the approved tool. Quiet workarounds are how organisations discover their real exposure eighteen months late.

What goes in the box

The single decision that matters most is what you paste. Get this right and almost everything else is recoverable.

For a public or unapproved tool, imagine reading the material aloud at a conference. If that would expose information, keep it out. For an approved business tool, use the organisation's classification and purpose-specific permissions instead; confidential material still needs explicit authorisation.

For anything the test does not settle quickly, the table below covers most of what actually turns up.

MATERIAL	VERDICT	WHY
Your own notes, drafts and agendas	Fine	Only if you are authorised to share it and it contains no protected information

MATERIAL	VERDICT	WHY
Material already publicly available	Fine	Public availability does not remove copyright, licence or personal-data restrictions
Internal documents, plans, reports	Approved tool only	Commercially sensitive even when unmarked
Customer or supplier correspondence	Approved purpose; minimise first	Removing obvious identifiers may not anonymise the material; retain only what the approved task needs
Staff records, pay, performance, health	Never, without a specific approval	Employment records contain personal data; health data is special category data
Signed contracts and anything under NDA	Specific approval required	Keep it out unless the information owner has specifically confirmed that the contract and approved tool permit this use
Security secrets such as passwords and private keys	Never	Do not paste secrets. Handle non-secret architecture information under the approved security process
Anything you were told in confidence by a colleague	Never	Approval covers the organisation, not the trust

Stripping properly

Removing a name is not anonymising. If the remaining detail identifies one obvious person to anyone who knows the situation, you have changed the formatting and nothing else. Take out the role, the site, the dates and the distinguishing circumstance, or reduce the material to the question you actually need answered.

A Useful Reflex

Before pasting anything longer than a paragraph, ask what the worst realistic consequence would be if the text appeared in a search result next year.

For most material the answer is nothing at all, and you should stop worrying and get on with it. The reflex exists to catch the handful of cases where the answer is not nothing.

What to ask, and what to take elsewhere

These tools are unusually good at some things and confidently poor at others, and the boundary is not where most people assume it is.

Where it earns its place

- Reshaping material that already exists: summarising, restructuring, changing the register for a different reader.
- First drafts of routine writing, where you were going to rewrite it anyway and the blank page is the expensive part.
- Explaining something unfamiliar at whatever level you ask for, then answering the follow-up you were embarrassed to ask a colleague.
- Mechanical work at volume: reformatting, extracting, tidying a list somebody sent badly.
- Arguing with your own thinking. Asking what is weak about a plan you wrote is one of the genuinely underused moves.

Where it is a poor choice

- Anything where you cannot check the answer and the answer matters. Confidence is not evidence, and you will not be able to tell the difference.
- Current facts, figures, prices and law, unless the tool retrieves and cites a source you then open yourself.
- Advice that belongs to a qualified person: legal, medical, financial, safety-critical engineering. Use it to prepare your questions, not to replace the answer.
- Anything requiring judgement about a specific individual. That is the next section.

Decisions that stay with a person

This template sets a policy default: consequential decisions stay with an accountable person. Specific legal requirements depend on the data, use and jurisdiction; they need their own assessment.

- Hiring, promotion, discipline and dismissal. A person makes the decision, records their reasoning, and can explain it without referring to a tool.
- Anything with a legal or similarly significant effect on someone: credit, eligibility, access to a service, allocation of work that affects pay.
- Safety decisions, and anything where being wrong hurts somebody.
- Communications that commit the organisation: contractual terms, regulatory responses, public statements.

Where This Comes From

The Data (Use and Access) Act 2025 changes the UK framework for significant automated decisions. Safeguards include information, human intervention and an opportunity to contest the decision; special category data remains subject to additional restrictions. Check the current ICO guidance for the proposed use.

The European Commission reports that, following the AI Omnibus, Annex III high-risk rules apply from 2 December 2027 and Annex I product-related high-risk rules from 2 August 2028. Employment uses can fall within the high-risk categories; classification depends on the particular system and use.

Article 50 transparency obligations apply from 2 August 2026. Duties differ between providers and deployers and by the kind of content or interaction; assess the relevant obligation rather than treating all AI-assisted work alike.

Reading the output

Fluency and accuracy are unrelated. The text is generated to read well, and it reads exactly as well when it is wrong.

This is the hardest habit to build, because everything about a polished paragraph tells your brain it has been checked. Nothing has been checked. What you are looking at is a plausible continuation, assembled at speed, with no internal sense of which parts it is sure about.

Four ways it goes wrong

WHAT YOU SEE	WHAT IS HAPPENING	WHAT TO DO
A specific detail that turns out not to exist	It filled a gap rather than leaving one	Check any specific claim you did not already know
A summary missing the thing everyone remembers	That thing was never written down anywhere it could read	Say what you expected. The gap is often the finding
Numbers that look right and do not reconcile	Arithmetic performed on a misread column	Total one column yourself before it goes anywhere
Guidance that was true last year	Trained on a world that has since moved	Ask for the source, then open the source

Checking in proportion

Verifying everything to the same standard is unworkable, so people quietly stop verifying anything. Scale the effort to what being wrong would cost.

- Low stakes, easily corrected, stays internal: read it once for sense and move on.
- Goes to a customer or a colleague who will act on it: check every factual claim and every number.
- Money, legal position, safety, or anyone's employment: open the sources yourself, and have a second person read it.
- You cannot tell which of those it is: treat it as the higher one until you can.

A Question Worth Asking The Tool

“Which parts of that came from the document I gave you, and which did not?”

It is not a guarantee, and it can be wrong about its own working. It is still one of the fastest ways to find the sentence that deserves your attention.

Bias, in the places it bites

Bias in these systems is not usually dramatic. It is a steady tilt in ordinary language, which is what makes it easy to ship without noticing.

The models learn from an enormous quantity of human writing, and human writing carries every assumption its authors held. Some of that is scrubbed out during training and some is not. What survives tends to show up as defaults: the seniority a role implies, the pronoun attached to a profession, whose English is treated as correct, which examples come to hand first.

Where it does real damage

- Recruitment. Job adverts written in language that narrows the field, and any sifting or ranking of candidates, which is exactly the use the regulators are most focused on.
- Performance and progression. Feedback rewritten by a tool that softens for some names and sharpens for others.
- Customer treatment. Segmenting, prioritising or scoring people, where a small tilt becomes a policy applied at scale.
- Anything about people in a group you are not part of. You are least likely to notice the tilt that does not point at you.

A short review routine

For anything touching people, four questions catch most of what matters and take about a minute.

- Would this read the same if the name, gender, age or nationality were different? Try it. Actually change one and rerun it.
- Has it introduced a requirement, a tone or an assumption I did not ask for?
- If this were applied to a thousand people rather than one, who would come off worst?
- Could I explain this wording to the person it describes, with them in the room?

Two Things That Are Not Solutions

Asking the tool whether its answer is biased. It will produce a reassuring paragraph, generated by the same process that produced the original.

Assuming that removing names removes bias. Writing style, school, postcode and career gaps all carry the same signals, and a system with no names at all can still sort people the same way.

Dos and don'ts

Everything above, compressed. If your organisation adapts this guide into a policy, this is the part people will actually read.

Do

- Use the approved tool for anything that touches your organisation, its customers or its people.
- Paste the real material when the tool is approved for it. Thin context produces generic output, and generic output wastes the time you were saving.
- Say what good looks like: the audience, the length, the tone, the format.
- Ask where a claim came from, and open the thing it points at.
- Read every line before it leaves you, including the ones that look boring.
- Tell someone when the output was wrong in a way that mattered. That is how the guidance improves.

Do not

- Paste organisational material into a tool nobody approved, however good it is and however quick the question.
- Put staff, health, pay or disciplinary information into any tool without a specific approval for that purpose.
- Let a tool make a decision about a person. Use it to prepare, then decide yourself and record why.
- Ship figures, quotations, legal wording or citations you have not opened.
- Turn off or work around a control because it is inconvenient. Raise it instead, and expect it to be looked at.
- Assume approval is permanent. Tools change terms, and what was fine last year may not be this year.

Who owns the output

Automation moves the typing. It does not move the responsibility, and it never has.

Whatever produced the draft, the person who sends it is accountable for it. That is not a policy position invented for AI; it is the same rule that has always applied to work delegated to a junior colleague, an agency or a template.

Disclosure

There is no universal rule about saying that AI was involved, and pretending otherwise helps nobody. What follows is a reasonable default that most organisations converge on.

- Ordinary drafting help, heavily edited by you: follow any applicable professional, contractual, legal and organisational disclosure rules.

- Substantially generated material that goes out largely as produced: say so, particularly to customers and in anything published.
- Analysis, research or recommendations others will rely on: name the tool and the method, because the reader needs it to weigh the conclusion.
- Synthetic images, audio or video: label them as a house policy. Where the EU AI Act applies, assess the specific Article 50 provider or deployer duty and its exceptions.

Records worth keeping

For anything consequential, keep enough to reconstruct how the output came about: the tool, the material you gave it, and what you changed afterwards. Six months later, when somebody asks how a figure was arrived at, this is the difference between an answer and an apology.

When it goes wrong

It will, occasionally. Speed of response matters more than the size of the mistake.

- Stop using the output. Do not send the corrected version until you understand what was wrong with the first one.
- Work out what actually left the building: who received it, what it said, and whether they have acted on it.
- If organisational or personal information went into a non-approved tool, report it the same way you would report any other data incident, immediately and without editing the story. Statutory reporting deadlines are short and start running from awareness.
- Correct it visibly with the people affected. A quiet replacement is usually discovered and always looks worse than the original error.
- Write down what happened in a sentence, and pass it to whoever maintains this guidance.

The Cultural Bit

An organisation that treats every AI mistake as a disciplinary matter will not hear about the next one until a customer does.

The behaviour worth discouraging is concealment, not error.

What this is built on

This is ITC's practical interpretation of the sources below, with policy defaults for an organisation to adapt. It is not a statement that every recommendation is a legal requirement.

SOURCE	WHAT IT CONTRIBUTES
NIST AI Risk Management Framework	The idea that risk is managed continuously through governing, mapping, measuring and managing, rather than signed off once
ISO/IEC 42001	A certifiable management system for AI, which is where an organisation goes when customers start asking for evidence
ISO/IEC 23894	Risk management guidance specific to AI, sitting alongside the general ISO 31000 approach
EU AI Act (Regulation 2024/1689)	Risk classification and use-specific transparency obligations; check current implementation guidance
UK GDPR and ICO guidance	Lawful basis, data minimisation and safeguards for significant automated decisions, including the DUAA changes
Sector regulators	Additional expectations in financial services, healthcare, legal practice and safety-critical engineering

Check the linked primary sources before adopting this as an internal policy, including the rules that apply in your jurisdiction.

Adapting this into your own policy

This document is deliberately publishable as it stands. To turn it into internal policy, five things need your organisation's answers rather than mine.

- Name the approved tools, explicitly, with the version or tier. "The approved AI tool" is not a list and people cannot follow it.
- Say who decides. Name the role that approves a new tool and the route for asking, so the answer to a request is a process rather than silence.
- Set the classification language to match your own. If you already have information handling tiers, map section 3 onto them instead of introducing a second scheme.
- State the reporting route for section 9 in one line, with the actual channel people should use.
- Give it a review date. Twelve months is too long for this subject; six is realistic.

Keep the tone. A policy that reads like a warning notice gets skimmed once and ignored, and the behaviour it was written to change carries on unobserved.

The one-page card

Everything that matters, on a single side. If your organisation adopts this guide, this is the page to put on the wall.

Approved tool for anything of ours

Company, customer or colleague material goes only into a tool the organisation approved.

Use other tools only for permitted general knowledge tasks with none of our detail in them.

Would you read it out at a conference?

Keep confidential material out of public or unapproved tools. Sensitive work needs specific approval for the data and purpose. Never paste passwords or private keys.

Fluent is not correct

Check any specific claim you would not have known. Open the source rather than trusting the citation.

Check in proportion to the cost of being wrong

Internal and reversible, read it once. Money, law, safety or someone's job, verify it and have a second reader.

People decisions stay with people

Hiring, discipline, performance, eligibility. Prepare with the tool, decide yourself, record the reasoning.

Test for the tilt

Change the name or the gender and run it again. If the output changes, so does your confidence in it.

Say something when it goes wrong

Report it the way you would any other incident, quickly and without tidying the story first.

Your name is on it

Whatever wrote the draft, you sent it. That has not changed and it is not going to.

Reuse

This guide is published as open guidance. Adapt it, rebrand it, cut it down and put it in front of people.

Check the current regulatory position and your organisation's requirements before publishing your own version.

Sources and further reading

ICO — Data Use and Access Act changes

<https://ico.org.uk/about-the-ico/what-we-do/legislation-we-cover/data-use-and-access-act-2025/the-data-use-and-access-act-2025-duaa-summary-of-the-changes/data-protection/>

European Commission — AI Omnibus

<https://digital-strategy.ec.europa.eu/en/news/ai-omnibus-enters-force>

European Commission — transparency guidance

<https://digital-strategy.ec.europa.eu/en/library/guidelines-transparency-obligations-providers-and-deployers-ai-systems>

NIST AI Risk Management Framework

<https://www.nist.gov/itl/ai-risk-management-framework>

Reuse this guide

Open guidance from IAMTHEITCONSULTANT. Adapt, rebrand and shorten it for your organisation. Confirm the current rules and obtain internal approval before adopting it as policy.