BEFORE YOU PASTE THAT INTO CHATGPT…

The 3-question test your whole team should run before putting company information into any AI tool — plus a starter policy you can adopt this week

Somewhere in your company, right now, someone is pasting something into a free AI chatbot. A customer list to "clean up the formatting." A contract to "summarize the key points." Pricing to "make this email sound better."

They're not being careless. They're being efficient. Nobody ever told them where that text goes.

WHERE PASTED TEXT ACTUALLY GOES

When an employee uses a **free or personal AI account**, the text they paste typically:

- Leaves your company's systems entirely
- Is processed under *consumer* terms of service — which in many cases permit the provider to use it to improve their models
- May be retained, reviewed, or logged in ways you can't audit and can't delete
- Is now outside every confidentiality control you've built

When a company uses AI under **enterprise terms**, properly configured, the picture is different: no training on your data, defined retention, contractual confidentiality. The technology is similar. The terms are not. **The risk isn't AI — it's the account it runs through.**

THE 3-QUESTION TEST (TEACH THIS TO EVERYONE)

Before pasting anything into an AI tool, ask:

1. Would I email this to a stranger? If no — client names, employee information, pricing, financials, anything under NDA — stop.

2. Is this a company-approved AI account, or a personal one? Personal accounts and free tiers are where confidentiality goes to die. If you're not sure which one you're using, assume personal.

3. If this text showed up somewhere public next year, could I explain how it got there? With a governed company tool: yes, there's a log. With a personal chatbot account: you couldn't even prove what was pasted.

If any answer makes you hesitate, don't paste. Ask instead: "What's our approved way to do this?" If your company doesn't have an answer, that's the real finding.

A STARTER AI-USE POLICY (STEAL THIS)

Most companies don't need a 40-page AI policy. They need five rules people can actually remember:

1. **Company information goes only into company-approved AI tools.** Personal and free accounts are for personal questions.
2. **Client-confidential material, employee data, and financials never go into a chatbot** — even an approved one — unless the tool runs *inside* our systems under our controls.

3. **AI output is a draft, not a decision.** A human reviews anything that goes to a customer, a regulator, or a court.
4. **If AI takes an action** (sends, files, orders, pays), **a person approves it first** — no exceptions until we've explicitly decided otherwise, in writing, per task.
5. **When in doubt, ask.** Nobody gets in trouble for asking. People get in trouble for pasting.

THE LONGER-TERM FIX

The paste problem exists because the AI lives *outside* your business and your data has to travel to it. The durable fix inverts that: AI that works *inside* your own systems — reading only the data you've permitted, under your access rules, with every action logged — so nobody has to carry your confidential information out the front door to get help with it.

That's not a product category you buy off the shelf. It's how modern custom software is built.