AI Use Policy

[Company Name]

AI arrived faster than anyone expected. Tools like ChatGPT, Copilot, and Gemini can draft a letter, summarize a long document, or work through a spreadsheet in seconds. Many of us are already using them, at home and at work. That is a good thing, and we are not trying to stop it.

Here is the part most people have never been told. These tools work by taking whatever you type or upload and sending it somewhere else to be processed. **Where that somewhere is, who can read it afterward, and how long it is kept depends entirely on which tool you use and how you signed in.**

This document is not a list of things you cannot do. It is a map. It shows which kinds of work are fine to hand to any AI tool, which need a company account, and which need to stay somewhere we control. Most of what you do every day falls in the first two groups.

Read it once — about ten minutes — and keep it somewhere you can find it again.

1 Who this applies to

Everyone at [Company]: employees, contractors, and temporary staff. It covers any AI tool used for any company purpose, on any device, whether the device belongs to you or to [Company].

2 What we mean by company information

Company information is anything you know or hold because you work here. That includes information about our customers, patients, vendors, partners, and co-workers — not only information about the company itself.

A simple test: if you would not have it if you did not work here, it is company information.

3 Your personal account is not the same as a company account

Think of a personal AI account the way you would think of a personal phone with company email on it and no controls. [Company] cannot see it, cannot manage it, and cannot switch it off when you leave.

- [Company] provides **company AI accounts, with a separate login for each person.** A company account comes with a written contract that stops the vendor from using our information to train their AI, and it lets [Company] manage the account. A free personal account does neither of those things.
- Anything above Level 1 happens in your company account.
- Your own personal AI account is fine for Level 1 work only.
- Company AI accounts are switched off on your last day, the same as your email.

4 The four levels

Every piece of information you might hand to an AI tool sits at one of four levels. The higher the level, the more careful we have to be about where it goes.

LEVEL 1

Public

Nothing about our company in it.

For example General questions, public facts, how something works, help with wording or formatting, practice and learning, anything [Company] has already published.

Where you may use it Any AI tool at all, including your own personal account.

LEVEL 2

Internal

Company business, nothing sensitive.

For example Meeting notes with nothing confidential in them, project plans, job descriptions, quotes from suppliers, how-we-do-it documents, marketing copy about our own services, everyday correspondence.

Where you may use it Your company AI account only — see section 5. Never a personal login.

LEVEL 3

Confidential

Would hurt us if it got out.

For example Financial statements and budgets, customer and patient lists, sales figures and deals in progress, forecasts, our pricing and profit margins, notes from confidential meetings, contracts, staff and HR files, strategy documents, source code.

Where you may use it A private AI system that we control. An ordinary internet AI tool may only be used here if it has been approved for Level 3 in writing — see the addendum at the back.

LEVEL 4

Restricted

Protected by law, or by a promise we made.

For example Patient health information, Social Security numbers and similar personal identifiers, credit card and bank details, anything discussed confidentially with a lawyer, anything covered by a non-disclosure agreement, trade secrets, passwords. Also anything governed by rules such as HIPAA (health), PCI (card payments), FINRA or SEC (investments), GLBA (banking privacy), or ITAR (defense exports).

Where you may use it A private AI system that we control, or not used with AI at all. Never an ordinary internet AI tool, no matter whose account it is.

> **Not sure which level something is?** Treat it as the higher one and ask **[Name, role]** at **[email]**. Nobody here gets in trouble for asking. Guessing is the thing this policy exists to prevent.

5 Which tools are approved

Only the tools listed below may be used for Level 2 and above. If a tool is not on this list, it is Level 1 only until someone adds it.

Tool	Approved up to	Which account	Ask for access from
Example: Copilot	*Level 2*	*Company*	*IT*

6 Things we never use AI for

These apply no matter which tool you use or what level the information is:

- Giving legal, medical, or financial advice to anyone outside [Company].
- Deciding, or helping decide, who gets hired, promoted, disciplined, or let go.
- Producing anything we file with a regulator or certify as accurate, unless a person has reviewed it and signed off.
- Pretending to be a real person, or passing AI-written work off as someone else's.
- Writing or changing code so that it gets around a security check.
- Handling information belonging to a customer or partner whose contract says we may not.

7 Six habits that apply everywhere

- Never type or paste a password, a login code from your authenticator app, or an API key (the password one piece of software uses to talk to another) into an AI tool.
- Treat everything AI gives you as a first draft. Check it before you send it, sign it, bill from it, or act on it.
- Do not let an AI tool connect to your company email, files, or calendar, and do not switch on new AI features in software we already use, without asking first.
- Assume anything you type into a public AI tool could be read by someone else one day.
- If you think you have put something in the wrong place, tell [Name] the same day. Mistakes caught early are almost always fixable.
- None of this changes because you are busy, working from home, or on your phone.

This policy follows the U.S. government's guidance on managing AI risk — the NIST AI Risk Management Framework, at nist.gov/itl/ai-risk-management-framework. We review it at least once a year, and any time a new AI tool is approved.

Addendum — a closer look at Microsoft Copilot

Copilot is the AI tool most of us already have, so it is worth understanding properly. Everything on this page comes from Microsoft's own published documentation. It is here so that a decision about which levels Copilot is approved for gets made on facts instead of assumption.

This applies to Copilot when you sign in with the company Microsoft account IT set up for you — signing in with a personal Microsoft account is a different product with different rules. Sources, all from Microsoft Learn: Data, privacy and security for Microsoft 365 Copilot; Enterprise data protection in Microsoft 365 Copilot; Copilot Chat privacy and protections; and the default policy for spotting sensitive information in Copilot. Microsoft updates these from time to time — check the current versions before relying on this page.

What Microsoft promises

- What you type, and what Copilot answers, stays inside Microsoft's own systems.
- It is not used to train the AI models.
- It is not handed to OpenAI to use.
- Copilot only shows you files you already had permission to open.
- It follows the confidentiality labels, record-keeping rules, and activity logs [Company] has already set up.

What it does not do

- **Your conversations are saved.** They are stored in the company email system and can be pulled up later by IT, HR, or lawyers. They are company records, not private notes.
- **There is no set expiry.** How long they are kept depends on settings [Company] chooses. If a lawsuit or investigation begins, they can be locked in place indefinitely.
- **It does not fix file permissions.** If a folder was shared more widely than it should have been, Copilot will find what is in it — much faster than a person would.
- **The standard Microsoft safety net starts switched off.** Out of the box, the tool that spots sensitive information in what you type only makes a note of it. It does not stop it, until an administrator turns that on.
- **It may search the web.** If that feature is on, a few words drawn from what you typed get sent to Bing.
- **Outside AI companies are involved.** Microsoft lists both OpenAI and Anthropic as suppliers of the AI models behind Copilot.

Five questions to get answered in writing

Before any ordinary internet AI tool is approved for Level 3 work, [Company] should have written answers to these:

- Exactly which product and which licences are approved, and who checked?
- Which levels of information are allowed, and which are not?
- Is the sensitive-information safety net actually switched on, or only taking notes?
- How long are conversations kept, and who can go back and read them?
- Are web search, add-ons, and connections to other software allowed when confidential information is involved?

> None of this makes Copilot unsafe, and none of it is a criticism of Microsoft. It simply means Copilot is **a well-run company filing cabinet, not a private safe.** Whether that is good enough for a particular kind of information is a decision for [Company] and its advisers, not something to assume. Until the five questions above have written answers, this policy approves Copilot up to Level 2.