

THE AI EMPLOYEE HANDBOOK

A fill-in-the-blanks governance template for any AI you put to work — write one page, stay in charge

Companies have employee handbooks because working relationships go wrong in predictable ways — and writing the rules down *before* the problem is what keeps small issues small.

AI in your business deserves the same one page. Fill this out for **each** AI agent or AI feature you adopt (yes, including the ones your existing software vendors just switched on). If you can't fill out a section, that's not a paperwork gap — that's a control you don't have.

1. Position description

This AI's job is to: _____ It explicitly does NOT: _____
_____ It works on behalf of (department/owner): _____

If the job can't be written in two sentences, the job is too big. Split it.

2. Access granted

Systems and data it can read: _____ Records it can create or change: _____ Everything else: no access — not "unused access." No access.

Rule of thumb: the AI gets the same permissions as the most junior person who could hold this job — never more.

3. Supervision — the approval matrix

ACTION	CAN DO ALONE?	REQUIRES SIGN-OFF FROM
Draft (documents, emails, entries)	Yes	—
Send anything outside the company	No	_____
Move money / issue credit / place orders	No	_____
Change a record a regulator or auditor might read	No	_____
Delete anything	No	_____

Start strict. Loosening a rule later is a decision; discovering it was never there is an incident.

4. Recordkeeping

Every action this AI takes is logged with **what, when, and under whose authority**, and the log is kept for _____ (months/years), reviewable by _____.

5. Data rules

Data it may use: _____ Data that never reaches it:
_____ Our data **is / is not** used to train the provider's models
(get this in writing — enterprise terms, not a sales email).

6. Performance review

Every _____ **weeks**, _____ samples the log and the output and answers: Was the work right? Did anything skip an approval? Expand the role, keep it, or shrink it?

7. Termination procedure

To shut this AI off immediately, _____ does _____, and it takes effect **instantly / within _____ minutes**, including tasks already in progress.

Test it once before you rely on it. An off switch you've never flipped is a rumor.

WHY ONE PAGE IS ENOUGH

Notice what this handbook doesn't ask: how the model works, what a token is, which vendor's AI is "smartest." None of that is your job to know — any more than you needed to understand your accountant's calculator.

What *is* your job is exactly what this page covers: **scope, access, supervision, records, and the exit**. Those five things are how you manage people, and they're how you manage AI. A vendor who can satisfy this page is offering you an employee. A vendor who bristles at it is asking for the master key.

Slingr builds custom business software where this handbook isn't aspirational — it's the architecture. AI agents run under scoped permissions, every action is logged, approvals are wired in, and the off switch is real. slingr.io