

AI Job Description Pack

Complete this for one agent before the licence is renewed. It takes about an hour and it is the artefact a pilot needs to reach production.

What it does. It defines what a system is hired to do, what it is forbidden to do, where its authority stops, how its errors are captured, and whose judgement is on the line when it is wrong.

What it is not. It is not a compliance artefact and it does not establish a legal position under any regulation. It is a management document.

Who should complete it. The person who requested the tool, the person who owns the workflow it sits in, and the person whose name will end up in section 4. If that third person is not in the room, stop.

Agent or tool

Date

How to use this pack

One agent per pack. A pack covering "our AI tools" collectively is the problem it is meant to solve.

1 • Statement of work

Page 3. What it is hired to do, and the explicit prohibitions. The prohibitions matter more than the purpose.

2 • Authority line

Page 4. Where its authority stops, and the worst single unsupervised action it could take.

3 • Probation

Page 5. How errors are logged as types rather than quietly corrected.

4 • Consequence owner and placement

Page 6. The named person, and where the agent sits in the value stream.

5 • Fallback rehearsal

Page 7. What shipping looks like when the tool is unavailable, wrong or compromised.

A test before you begin. If nobody can state what the agent is forbidden to do, it has not been designed. It has been bought.

SECTION 1

Statement of work

Written as it would be for a contractor. Specific enough that someone else could tell whether it had been breached.

What this agent is hired to do

One bounded task inside one workflow. Not "help with engineering".

What it is explicitly forbidden to do

List the actions that would be a breach even if the output looked correct.

Expected outputs, and what "good" looks like

SECTION 2

Authority line

Least privilege, stated as actions rather than permissions.

ACTION	AGENT MAY DO ALONE	REQUIRES A HUMAN

Worst single unsupervised action

If it is wrong once, at the worst possible moment, what is the maximum damage—and what structurally prevents it?

What prevents it

SECTION 3

Probation

Correcting the output fixes today. Logging the error type fixes the system.

DATE	WHAT IT GOT WRONG	ERROR TYPE	CHANGE MADE TO PROMPT, CONTEXT OR SCOPE

Review at the end of probation. If the error log is empty, nobody is checking. If it is full and no change was ever made to prompt, context or scope, the feedback loop is decorative.

Consequence owner and placement

A natural person with competence, training and authority — the standard the EU AI Act sets for deployers of high-risk systems, and a reasonable bar for anything that acts.

Named consequence owner

Whose performance review is affected if this fails

Escalation path when the agent is uncertain or out of scope

Upstream supplier

Who or what provides the input, and who is responsible for its quality.

Downstream consumer

Who acts on the output, and whether they can tell it came from an agent.

Fallback rehearsal

Resilience is the ability to keep working when the primary system does not.

If this tool is unavailable tomorrow, how does the next release ship?

If its output is wrong and nobody notices for a week, what is the recovery?

Which tasks are deliberately kept manual to preserve human judgement?

The work that teaches a new joiner how the system behaves is not the same as friction to be removed.

Boundaries of this pack

The four elements and this pack are Beta Tester Life interpretations of published research on enterprise AI adoption and of ordinary job-design practice. No study establishes that organisations completing this reach production more often than those that do not. Quoting the EU AI Act here is a borrowed standard, not a legal obligation: most tools in a typical delivery organisation are not high-risk systems under the Act.

This pack is not a substitute for legal, regulatory, HR or professional advice.

Sources: MIT Project NANDA, *The GenAI Divide: State of AI in Business 2025*, July 2025; Gartner press release, 25 June 2025; EU AI Act, Article 26 and implementation timeline; DORA, *2025 State of AI-assisted Software Development Report*.