CloudWarrior

PLAYBOOK · AI AGENTS FOR BUSINESS

An AI agent. It does only what was agreed.

How we deploy an AI agent into one business process: with a scope, limits, a log and a set point where a person takes over.



01 · WHO IT'S FOR

Your AI assistant is "sometimes right."

Your team uses an AI chat, or already has an assistant that answers customers, sorts tickets or reads documents. It usually works. But nobody knows what it will do in an unusual case, where its answer came from or what data it can access.

An agent without written limits is a risk that grows with its permissions. This playbook shows how to set those limits before the agent gets real data.

This is for you if

- the task is text-based: emails, tickets, documents,
- it repeats many times a week,
- a person can easily check the result,
- someone can take over the exceptions.

It's not the right time yet if

- the agent would make decisions with no oversight,
- there are no sample cases to test with,
- the process rules aren't written down anywhere,
- the goal is to train your own model.

What you get

- An agent with a clear scope: what it does and what it never touches.
- Limits: which systems it can access and what it can do without a person.
- A log of every step, so you can see where each answer came from.
- Handover to a person at an agreed point.

02 · HOW IT WORKS

Five steps from first call to handover.

01 **Call**

Free. You show us where AI works today, or where it should, and examples of cases it gets wrong.

02 **Diagnosis**

Paid. We write down the agent's scope, limits, data sources and acceptance criteria. The cost is credited toward an implementation ordered within 30 days.

03 **Build**

On your accounts. We connect the agent to the systems it needs with the least access possible. We show you the first cases live.

04 **Acceptance session**

On a copy of your data. We check normal cases, out-of-scope cases and the point of handover to a person. The result goes into the acceptance record.

05 **Handover**

The configuration, the log and a "what to do when the agent gets it wrong" guide go to you. You pay model fees directly to the provider.

Payment: 50% before start, 50% after the signed acceptance record. We set the timeline and price in the diagnosis.

03 · ACCEPTANCE

We also test what the agent must not do.

An example for an agent that sorts customer tickets. The data is illustrative.

#	CRITERION	HOW WE CHECK
1	Each ticket goes to the right category and person.	A set of sample tickets.
2	An out-of-scope question goes to a person, with no reply from the agent.	A deliberately unusual ticket.
3	The agent has no access to data it doesn't need.	A review of permissions.
4	Every decision has a log entry with its basis.	Reading the log after the test.
5	Someone outside the project turns the agent off using the guide.	A run with the guide in hand.

What we need from you

- Written rules: what the agent is allowed to do.
- Sample cases.
- A person who takes over exceptions.
- One acceptance session.

What's not in scope

- Training your own model.
- Model fees.
- Content the client is responsible for.

Next step

Collect three examples where your current assistant got it wrong, or describe the task in two sentences. The initial analysis and audit are free; we propose scope, cost and timeline for implementation afterwards.

Describe your task: cloudwarrior.io/ai