Mindlab Systems

ONE-PAGE GUIDE · FREE TO REUSE

PROVE: five steps to AI agents that pay off

Write down what success looks like and set the rules before anything is built. Works for agents on Microsoft Copilot Studio, Claude or ChatGPT.

P	R	O	V	E
Pick	Record	Own the guardrails	Validate	Evidence

P Pick

One process, one owner, one written target.

Choose a task your team repeats every week, name the person who owns it, and write down the result you want, for example “invoice triage from 3 days to 4 hours”.

You end with: A one-line goal with an owner

R Record

Measure how the work runs today.

Hours spent, wait time, errors and cost, from two weeks of tracking or the data you already have. Without a baseline, nobody can say later whether the AI helped.

You end with: A baseline card

O Own the guardrails

Set the rules before anything is built.

Who owns the agent, what it may read, what it may do, who approves its actions, and a monthly spend cap. The same rules apply whether it runs on Copilot, Claude or ChatGPT.

You end with: An Agent Fact Sheet

V Validate

Build small, then test with real questions.

Build the smallest agent that does the job on the right platform, test it with at least 20 real questions or cases, fix what it gets wrong, and only then release it.

You end with: A test set with a pass rate

E Evidence

Compare with the baseline 30 days after go-live.

Measure the same things as in step R. Then decide: scale it, change it or switch it off. Start the next process with what you learned.

You end with: A results report and a decision

Measure before you build.
The same rules for every AI tool.
Guardrails before agents.
Switch off what doesn’t pay.

© 2026 Mindlab Systems Inc. Share it freely with credit.