

AI Opportunity Map

Write up one candidate use case properly, or see what a Sprint delivers.

A one-to-two-page structure for deciding whether one piece of work is worth taking to production with AI. It is the deliverable at the end of a Sprint, and it works just as well as a self-assessment. Fill it in for one candidate at a time. Write it for the person who has to decide, in plain words.

How to use it

1. Pick one piece of work. Not "customer service": "answering order-status emails".
2. Fill in sections 01 and 02 before you think about technology. If you can't, you don't know the work well enough yet. Go and watch it.
3. Fill in the rest. Where you are guessing, write "assumed". Where you don't know, write "unknown". Both are fine. Hiding them is not.
4. Comparing several candidates? Score each finished Map with the use-case scoring rubric.

01 • The work

What is the job, who does it, how often, and what does "done" look like? One paragraph, no technology words.

Field	Answer
The job	
Who does it (roles, how many people)	
How often (per day or week)	
What "done" looks like	

02 • Where it hurts

Where do people lose time, re-enter data, wait for someone, search across documents, or rely on what one person knows? Be specific.

Friction	Minutes per instance	Instances per week	Who is affected	Measured, assumed or unknown

03 · The candidate future workflow

Describe the job again with the proposed change. Say what a person does, what the system does, and where a person checks the system's output before it counts.

04 · Recommendation

Pick one, and say why the alternatives lost.

Option	Chosen?	Why it won or lost
AI		
Simpler automation (rules, forms, integrations)		
Process change (who does what, when)		
No change		

05 · What it would need

Area	What it needs	State today
Data (what, where, how clean)		
Tools and integrations		
Who maintains it in production		

06 · Controls

Answer each one. "Not applicable" needs a reason.

- **Privacy and security.** What data does it touch? Who may see it? Where may it be processed?
- **Human review.** What must a person check before an output counts? Who?
- **Failure handling.** How is a wrong output caught, logged and escalated? What happens when the system is down?
- **Cost.** What does one use cost to run, roughly, and what is the monthly limit?

07 · Pilot scope and production exit criteria

Field	Answer
What the pilot must prove	
With whom, and for how long	
Explicitly out of scope	
Exit criteria: once met, it goes live	

Field	Answer
Stop criteria: if hit, it stops	
Who owns it after go-live	

08 · Success measures

The two or three numbers that would change the decision, and the baseline you need before starting.

Measure growth as well as time: quote or response turnaround, requests or orders handled per person, share of enquiries answered the same day.

Measure	Baseline today	Target	How it is measured

09 · Effort, risks and the next decision

Field	Answer
Rough effort (people, weeks)	
Biggest risks and how they are handled	
The one decision the sponsor makes next	
Decision owner and date	

Worked example (fictional)

A fictional 60-person industrial supplies distributor, for illustration only.

- **01 The work.** Three customer-service staff answer "where is my order?" emails, about 40 a day. Done means the customer has a correct delivery date.
- **02 Where it hurts.** Each answer takes about 6 minutes: look up the order, check the carrier site, write the reply. About 200 a week. Measured over one week.
- **04 Recommendation.** Simpler automation first: a nightly carrier-status import and a reply template cut most of the lookup. AI drafting of the unusual cases is a phase two candidate. AI alone lost because the hard part was data access, not writing.
- **07 Exit criteria.** Replies correct in 98% of a 100-email test set, median handling under 2 minutes for four weeks, no customer-facing reply sent without a person approving it.