My first AI MVP

Name:

Candidate:

Opportunity Definition

1 What is your main goal for using AI?

Keep asking "why" until you are sure you have arrived at the ultimate reason for making an investment into using AI. For example *"so it can do repetitive tasks we used to do manually"* may be a great idea, but it is not the ultimate reason why. The ultimate reason may be *"so we can deliver more of X in the same timeframe, allowing us to increase customer satisfaction"*.

2 Where is your bottleneck? Where is work piling up?

AI is great at generating content. However, if that content then clogs up your bottleneck, you are creating extra noise and backlog. This usually results in reduced efficiency, and can be quite dramatic. Instead ask yourself, what tasks are preventing us most from achieving our goal from step 1 (eg. slowing us down). That is your bottleneck, and it holds everything else back. How can you best relieve your bottleneck? That is where a first instalment of AI may help you most.

3 What can AI do that wasn't practical with humans?

Using AI does not always mean to automate tasks that were previously done manually. AI can do things that simply weren't practical for humans, for example producing many versions of something, or analyzing very large amounts of data. How can AI help achieve your goal from step 1 in ways that weren't previously possible?

MVP Scope

4 Mitigating Data-Leak Risks



Mark the corners your AI agents will touch.
A single agent should be able to do at most 2, never 3 of the corners
If one agent passes information to another agent so that together they can do all 3, you must sterilize/audit the data in between agents

Data can leak inside your organization, or to the outside. Anything that is entered into a computer should be treated as suspect. Unless you are willing to use only paper, there is a possibility your data might leak. This may not only be embarrassing and/or costly, it could also become a legal liability to your employees and customers.

Since staying on paper is not realistic for most organizations, AI adoption isn't a question of "if", it's a question of "how do we minimize risk". Several strategies are available, such as the data management triangle. Write down which strategies you plan to use.

- [ ] Data management triangle
- [ ] Opt out of data-retention, data-sharing, and training agreements
- [ ] Fence off data so you only share what is necessary
- [ ] Weigh the benefits against the risks. Ask: "What if this data got leaked?"
- [ ] Send data only to/through vendors you trust. Check for certifications
- [ ] Consider geographical boundaries. Look for local inference/ownership

5 What is delegated to AI vs remains with humans?

Techniques exist to increase the quality of AI's outputs (eg. adversarial agents, multi-agent chains/pipelines).

But some decisions are too risky to be made by machines. Which decisions must still be made by humans in your workflow?

How can AI make the jobs that remain with humans easier? (eg. templates)

What will humans do with the time AI freed up?

6 Your signal to see if it improved things, and whether to adapt/expand

When you look back to today in 2 weeks or months, how will you measure if you have moved closer to your goal from Step 1? If it didn't have the desired effect, what will you do? If it did, how can you get more from it, what will you do next?