

AI Use Case Canvas

The AI Use Case Canvas is a one-page template you fill in before you build anything, to decide whether AI is the right tool for a specific problem and whether this is the use case worth backing first. It maps the problem, the current cost, the data you have, the human review point, the way you will measure success, and the build-or-buy call. It earns its place when an idea sounds exciting but nobody has yet written down what would make it succeed.

Frame the problem

Problem statement

What to capture: The real pain in one or two sentences: who feels it, how often, what it costs. Write the problem, not the AI.

Example entry: 40% of ~120 weekly tickets are the same repeat questions that sit unanswered for hours.

Current cost of doing nothing

What to capture: The status-quo baseline in real units: hours, dollars, error rate, or wait time.

Example entry: Two engineers lose ~6 hours each per week; median first response is 4 hours.

Why AI (and not a simpler tool)

What to capture: The reason a rule, script, report, or process fix will not do the job as well or cheaper.

Example entry: Static FAQ is ignored; questions are phrased a hundred ways, so keyword routing misses them.

Check feasibility

Data and inputs

What to capture: The data the use case needs, where it lives, who owns it, and whether it is accessible and clean today.

Example entry: 18 months of resolved tickets we own and can export; messy but labeled; redact names first.

Human in the loop

What to capture: The review or approval point, who staffs it, and what they keep authority over. Flag any unreviewed action.

Example entry: Assistant drafts and cites a doc; a support engineer approves before send; AI never auto-sends.

Build or buy

What to capture: The build-versus-buy call and the reason. Default to buy for generic capability.

Example entry: Buy: two vendors ship suggested-reply features that read our docs; building would take a quarter.

Decide and de-risk

Success signal and horizon

What to capture: The one measurable signal that means this worked, its target, and the date you expect movement.

Example entry: Median first response on top 5 question types drops from 4 hours to under 30 minutes in 8 weeks.

Risks and failure modes

What to capture: What a confident wrong answer breaks, who is hurt, and how you would catch it. Include privacy and bias.

Example entry: Wrong permissions advice misleads a user; mitigated by human approval and an escalate fallback.

Decision

What to capture: The call (back now, defer with a named condition, or drop) and the owner.

Example entry: Back now as a 2-week pilot on the top 5 questions; owner Priya; expand if the signal holds two weeks.
