LAB 6 Worksheet 2: Pilot metrics charter

Type into the boxes, or print this and write on it. To fill it in on a phone, with examples to start from, use the worksheets page: **skillrepos.github.io/tlsaccel**

What this is: a one-page plan for your pilot. It says who runs it, which one measurement should improve, what must not get worse, and when you would stop. Each box says what goes in it and gives an example from a made-up company. Boxes marked **Required** are the ones the lab asks for today; **Optional** ones are for back at work. Nothing has to be exact today: a best guess, marked "guess", is a good answer.

LAB 6, STEP 2

Pilot header

Use case (required)
The use case you picked as your winner on Worksheet 1.
Example: Triage inbound support tickets

Workflow owner (required)
The one person who will run the pilot day to day. One person, not a team: a job title, or a name (a made-up one is fine).
Example: Dana Ruiz, Support Operations

Sponsor (required)
The one person who pays for the pilot and can clear roadblocks, often the workflow owner's manager: a job title, or a name (a made-up one is fine).
Example: Priya Shah, who runs Customer Service

Pilot team (required)
The one group of people who will use the AI tool during the pilot, and about how many they are. No names needed.
Example: One support shift: 12 of the 45 agents, plus the helpdesk administrator

Start date (required)
The first day the pilot team uses the AI tool in real work.
Example: January 5, 2027

Day-90 date (required)
The last day of the pilot: 89 days after the start date (the start date is day 1). On the worksheets page this is worked out for you; change it if your pilot must end on another day.
Example: March 31, 2027 (this team chose the end of the quarter, a few days short of 90)

Buy / partner / build (optional)
How you will get the AI tool. Buy: a product that already does this job. Partner: a vendor fits a tool to your data and way of working. Build: your own engineers make it. Most first pilots buy.
Example: Buy: the add-on our helpdesk vendor sells

(c) 2026 Tech Skills Transformations and Brent C. Laster. All rights reserved.

LAB 6, STEPS 3–4

Primary metric (pick ONE)

Family	Example metrics
Time	hours per task, cycle time, backlog age
Cost	cost per ticket/document/case, contractor spend
Quality	error/rework rate, first-pass acceptance
Adoption	weekly active users, % of volume through the tool

My primary metric (required)
The one measurement that should improve if the pilot works. Say what is measured and in what unit (minutes, dollars, percent or a count).
Example: Median handling time per support ticket (minutes)

Baseline plan (required)
How you will find out what your primary metric is today, before AI: who looks it up, from where, and for which two weeks.
Example: The helpdesk administrator pulls handling times for the pilot shift from the helpdesk, for two weeks in December

Baseline value (required)
What your primary metric is today, before AI. Not sure? Write your best guess and add "guess": measuring it is the pilot's first job.
Example: 11 minutes (handling time per ticket today)

90-day target (required)
What you want your primary metric to be by day 90, in the same unit as the baseline value. A number, not "better".
Example: 8 minutes (down from 11)

LAB 6, STEP 5

Quality guardrail

The number that must NOT get worse while the primary metric improves:

Guardrail (required)
A second measurement, one that shows quality. It must not get worse while your primary metric improves.
Example: First-contact resolution (the share of tickets solved in one reply)

Current value (optional)
What your guardrail is today. A guess is fine if you mark it.
Example: 72% (of tickets are solved in one reply today)

(c) 2026 Tech Skills Transformations and Brent C. Laster. All rights reserved.

LAB 6, STEP 6

Definition of done

"By **[date]**, **[workflow]** handled with AI assistance shows **[metric]** moving from **[baseline]** to **[target]**, with **[guardrail]** maintained — or we stop and write down what we learned."

My definition of done (required)

One sentence built from the boxes above: the day-90 date, the use case, the primary metric with its baseline value and 90-day target, and the guardrail.

Example: By March 31, 2027, support-ticket triage handled with AI assistance shows median handling time per support ticket moving from 11 minutes to 8 minutes, with first-contact resolution held at or above 72% — or we stop and write down what we learned.

LAB 6, STEP 7

Stop rule

A stop rule is the result that would make you end the pilot early, decided before you start.

Stop check on day (required)

The day of the pilot, about halfway through, when you look at the primary metric and the guardrail and decide whether to carry on. A number from 1 to 89.

Example: 42

We stop early if we see (required)

What the primary metric or the guardrail would have to show on that day for you to stop the pilot early.

Example: Handling time is still 11 minutes or more, or fewer than 72% of tickets are solved in one reply for two weeks running

LAB 6, STEPS 8–9

Finance red-team findings

A red-team is a deliberate attack on your own plan, to find its weak spots before someone else does. In step 8 an AI reviewer plays a skeptic from finance and names three weak spots, each with a fix. In step 9 you write down the two that matter most.

Most gameable part and the fix (required)

"Gameable" means the primary metric could look better without the work getting better. Write the weak spot the AI reviewer found in your plan, and how you will close it.

Example: Agents can look faster by splitting one problem into two tickets. Fix: also count reopened tickets and tickets per customer problem.

(c) 2026 Tech Skills Transformations and Brent C. Laster. All rights reserved.

Attribution risk and the fix (required)
Would the primary metric have improved anyway, without the AI? Write why it might, and how you will tell the difference.
Example: Ticket volume dips every January, so handling times fall anyway. Fix: compare the pilot shift with a shift that does not use the tool.

LAB 6, STEP 10

What the pilot needs

For step 10 of Lab 6, after the AI review. What it will take to run the pilot for 90 days: money, people, and a limit on what the AI tool may spend. Rough answers are fine.

Money (required)
What the pilot will cost in money over the 90 days: licenses, usage fees, vendor help. A rough figure is fine.
Example: $18 per agent per month × 12 agents = $216 a month

People and time (required)
Who will work on the pilot, and about how many hours a week each.
Example: Dana, 4 hours a week; each pilot agent, one hour of training

AI running cost per task (optional)
What the AI tool charges to do the job once, for example one ticket or one document. "Not measured yet" is a fine answer today: it is your first question for the team.
Example: Not measured yet: first question for the vendor

Monthly cap (optional)
The most you will let the AI tool spend in one month.
Example: $500

Who gets the alert (optional)
The one person who is told when the month's spending gets close to the cap.
Example: Dana Ruiz

LAB 6, STEP 10

If it succeeds

Also for step 10 of Lab 6. What happens after day 90 if the pilot works: who keeps it going, and what it costs once everyone uses it.

Who runs it after day 90 (required)
Who keeps the AI tool running once the pilot is over.
Example: Support Operations, under Dana Ruiz

(c) 2026 Tech Skills Transformations and Brent C. Laster. All rights reserved.

Rough cost at full volume, and who pays (required)

What the AI tool will cost when everyone who should use it does, and whose budget covers it.

Example: $18 × 45 agents = $810 a month, about $9,700 a year, from the Customer Service budget

TAKE-HOME

Seven questions for the team proposing the pilot

Tick each one you already have a clear answer to. The unticked ones are your next conversation.

- [ ] 1. What business problem are we solving — and who has it?
- [ ] 2. Why does this need AI? Would simpler, rule-based automation do?
- [ ] 3. What data does it need — does it exist, is it accurate, may we use it, and do we agree what the numbers mean?
- [ ] 4. What happens when the AI is wrong — who notices, who fixes it?
- [ ] 5. How much can it do on its own — what needs a person to approve?
- [ ] 6. How will we know it worked — baseline, target, guardrail?
- [ ] 7. What will it cost to run and to scale — not just to pilot?

(c) 2026 Tech Skills Transformations and Brent C. Laster. All rights reserved.