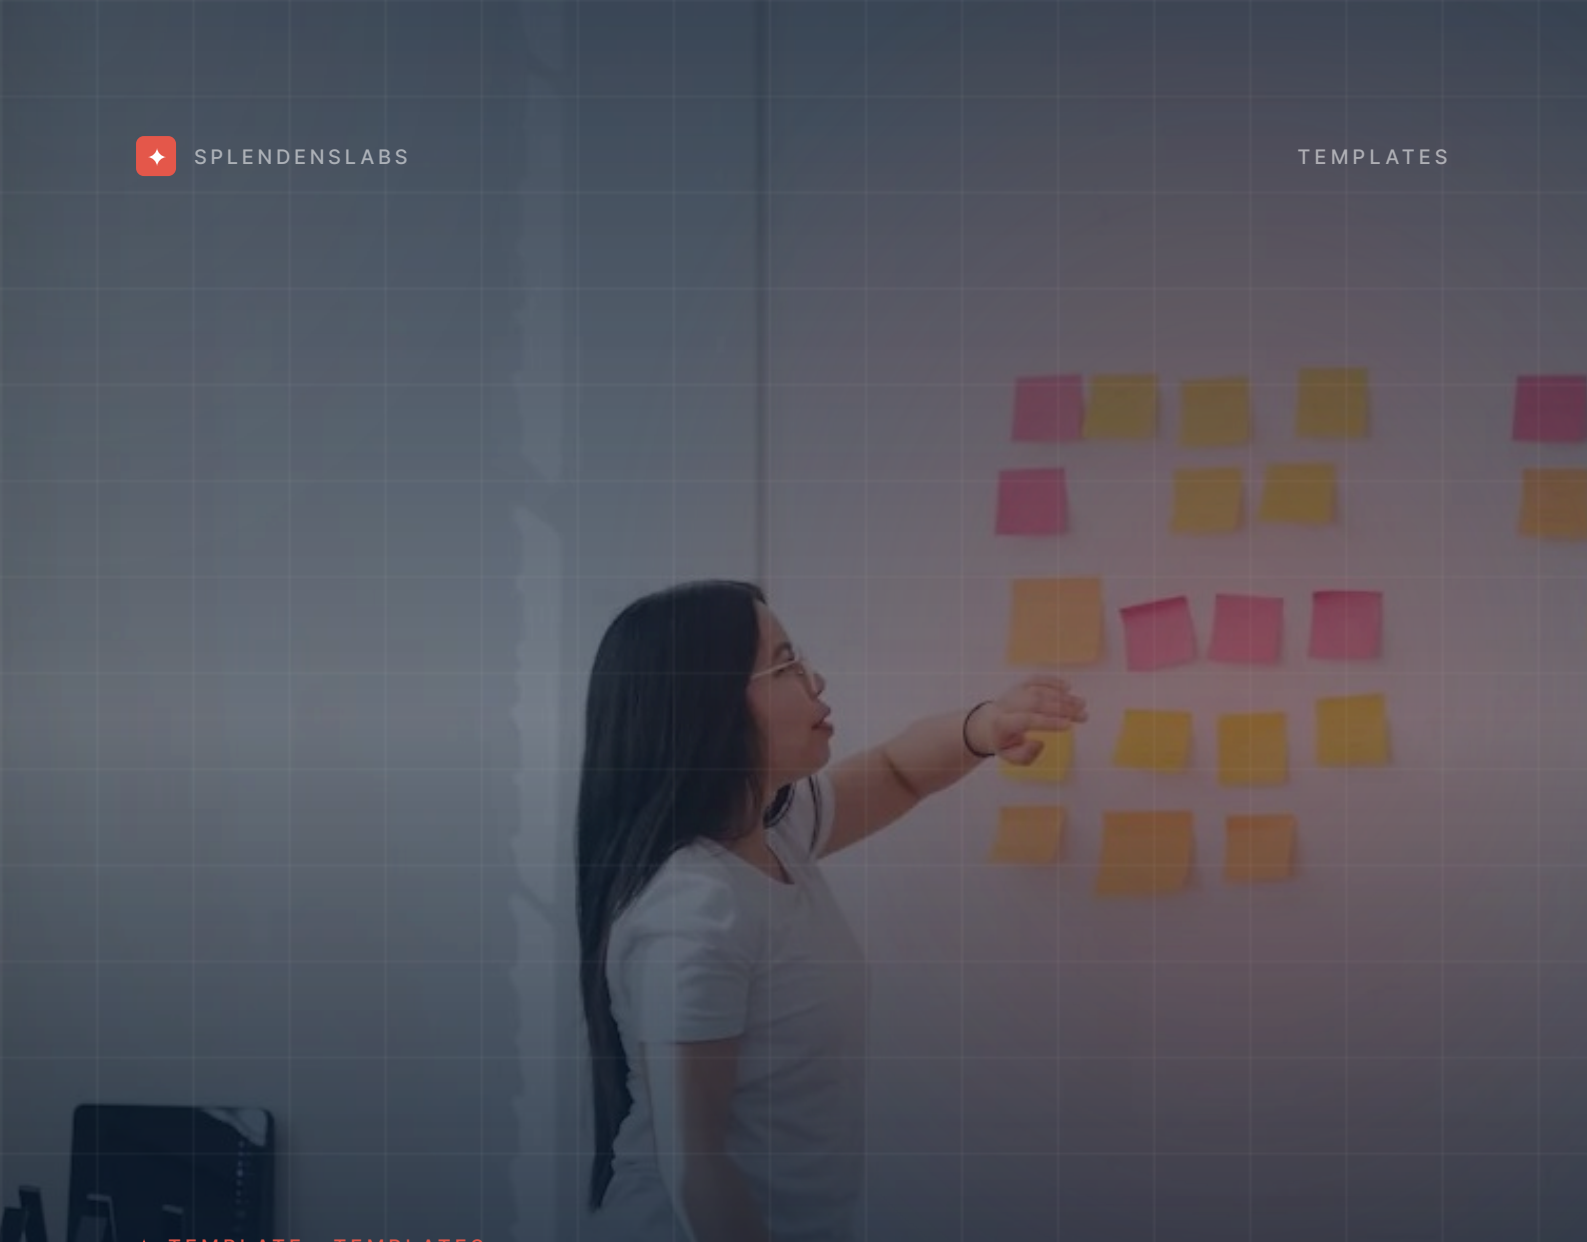


✦ TEMPLATE · TEMPLATES

AI Project Charter

A one-document charter that pins down the problem, the outcome, the data, the risks and who decides — before anyone writes a prompt.

FORMAT	LEVEL	READ TIME	UPDATED
Template	Beginner	Fill-in	Oct 2026

OVERVIEW

What this template covers

AI projects drift when nobody wrote down what success looks like, which data the system may touch, or who can say stop. This charter is the page we ask sponsors and delivery leads to complete together in the first week. It turns enthusiasm into a scoped, owned and measurable piece of work.

KEY TAKEAWAYS

- Name one business problem and one primary metric; everything else is secondary.
- Write the baseline down now — you cannot prove improvement without it.
- Agree what the system must never do before agreeing what it should do.
- Every decision on scope, data access and launch needs a named owner.
- Set a kill criterion up front so stopping is a decision, not a failure.

WHO IT'S FOR

- Executive sponsors
- Product owners and delivery leads
- AI / ML engineering leads
- Risk, legal and data owners

HOW TO USE IT

Fill this in during a 90-minute kickoff with the sponsor, product owner and engineering lead in the room; write "unknown" and an owner rather than leaving a gap. Every box is shown filled in with a worked example — a claims triage assistant for a fictional mid-size insurer — so you can see the level of detail expected. Replace the examples with your own and re-sign the charter whenever scope changes.

CONTENTS

01 Problem and outcome
02 Success measures
03 Scope and users
04 Data, constraints and risks
05 Roles, decisions and timeline
06 Review and sign-off

01 / PROBLEM

Problem and outcome

Start with the business problem, not the technology. If the problem would still matter without AI, you are on the right track.

Project name and one-line purpose

01

e.g. “Claims triage assistant — sort incoming claims by urgency and route them to the right team.”

WORKED EXAMPLE

Claims Triage Assistant — read each new motor and health claim, score its urgency, flag likely fraud signals and route it to the right claims team within minutes of arrival.

The problem today

02

Who is affected, how often, and what it costs in time, errors or missed revenue. Describe the current process in plain words.

WORKED EXAMPLE

- About 1,800 claims a week arrive by email, portal and branch scan; two coordinators read and route each one by hand
- A claim waits 1.5 working days on average before a handler sees it, and 1 in 8 is re-routed
- Urgent hospitalization claims sit in the same queue as minor windscreen claims

Desired outcome

03

What is different for users and the business when this works? Avoid “use AI to…”; describe the end state.

WORKED EXAMPLE

Every claim reaches the right team, with an urgency tag and a short summary, within 30 minutes of arrival. Coordinators spend their time on the ambiguous cases instead of sorting the obvious ones, and customers with urgent claims hear from a handler the same day.

Why now

04

The trigger: a growing backlog, a regulatory change, a customer commitment, a cost target.

WORKED EXAMPLE

Claim volume grew 35% after the new health product launched in April, and the regulator's turnaround expectations for health claims are being written into our service charter for January 2027. Hiring more coordinators only moves the bottleneck.

✦ 02 / SUCCESS

Success measures

One primary metric decides whether the project worked. Guardrail metrics make sure it did not succeed by breaking something else.

Primary metric and target

01

e.g. "Median time to first response on support tickets, from [baseline] to [target] within [period]."

WORKED EXAMPLE

Median time from claim arrival to assignment to the correct team: from **1.5 working days** to **under 2 hours** by the end of the pilot (31 Jan 2027).

Current baseline and how it was measured

02

The number today, its source system, and the date range. If there is no baseline, the first task is to measure one.

WORKED EXAMPLE

1.5 working days median, 4.1 days at p90 — measured from claim-system timestamps (received → first assigned) for 1 Jul–30 Sep 2026. Re-routing rate of 12% taken from the same export.

Guardrail metrics 03

Things that must not get worse: accuracy, complaint rate, escalation rate, cost per transaction, staff workload.

WORKED EXAMPLE

- Misrouting rate stays at or below today's 12%, target 5%
- No urgent health claim is tagged "routine" on the weekly audit sample
- Fraud-referral volume to the investigations team does not rise by more than 10%
- Model cost per claim stays under ₹6

Kill criterion 04

The condition under which we stop or pivot, e.g. "If pilot accuracy is below [threshold] after [n] weeks of tuning."

WORKED EXAMPLE

If routing accuracy on the 500-claim golden set is below 90% after six weeks of tuning, or any urgent health claim is missed in the pilot audit twice, we stop and revisit the approach.

Choosing the primary metric

METRIC TYPE	GOOD EXAMPLE	WEAK EXAMPLE
Efficiency	Handling time per case, measured in the ticketing system	"Staff will be more productive"
Quality	Share of answers rated correct by reviewers on a golden set	"Answers will be better"
Revenue	Conversion rate on assisted sessions vs. control	"Customers will like it"
Risk	Policy breaches caught before release	"It will be safer"

✦ 03 / SCOPE

Scope and users

Write the out-of-scope list with as much care as the in-scope list. It is the one that protects the timeline.

Primary users and their context

01

Role, how many, where they work (desk, field, mobile), and what tools they already use.

WORKED EXAMPLE

Six claims coordinators at the central operations hub (desktop, claims system plus shared mailbox), and 40 claims handlers across motor, health and property teams who receive the routed claims in their existing work queues.

In scope

02

The specific tasks, channels, languages and document types covered in this phase.

WORKED EXAMPLE

- New motor and health claims arriving by email and the customer portal
- English and Hindi claim forms and covering letters
- Attachments: claim forms, hospital discharge summaries, repair estimates (PDF and photos)
- Outputs: team assignment, urgency tag, three-line summary, fraud-signal flags

Out of scope

03

What we are deliberately not doing yet, e.g. “No autonomous actions on customer accounts in phase 1.”

WORKED EXAMPLE

- Approving, rejecting or settling any claim
- Property and commercial claims (phase 2)
- Branch-scanned handwritten forms until OCR quality is tested
- Any message sent directly to customers

Human in the loop

04

Where a person reviews, approves or can override the system's output — and what happens on disagreement.

WORKED EXAMPLE

Coordinators see the suggested team and urgency and confirm or change it with one click during the pilot. Fraud flags only add a note for the handler; they never block or delay a claim. If the coordinator disagrees, their choice wins and the case is logged for weekly review.

✦ 04 / DATA AND RISK

Data, constraints and risks

Most AI projects slip on data access, not models. Surface every dependency and constraint now.

Data sources and owners

01

System, data owner, access method (API, export, connector), and whether personal or sensitive data is involved.

WORKED EXAMPLE

- Claims system — Head of Claims Operations; existing API
- Claims mailbox and portal uploads — IT Service Manager; mailbox connector
- Policy admin system — Head of Underwriting; read-only replica
- Personal and health data involved: DPO review before access

Constraints

02

Data residency, regulatory obligations (e.g. India DPDP Act 2023), approved model providers, budget ceiling, latency needs.

WORKED EXAMPLE

India-region hosting and approved model providers only; health data handled under the India DPDP Act 2023 and insurer outsourcing rules. Build budget ceiling ₹45 lakh; routing within 5 minutes.

What the system must never do

03

e.g. “Never give dosage advice”, “Never reveal another customer’s data”, “Never send email without approval.”

WORKED EXAMPLE

- Never approve, reject or change the value of a claim
- Never contact a customer, hospital or garage
- Never put one customer’s details in another customer’s summary

Risk register

RISK	LIKELIHOOD (L/M/H)	IMPACT (L/M/H)	MITIGATION	OWNER
Health data sent to an unapproved provider	L	H	Gateway allows approved providers only	A. Das
Urgent claim tagged as routine	M	H	Coordinator confirms every tag; weekly audit	R. Menon
Weak accuracy on Hindi or scanned forms	M	M	Own golden-set slice; out of scope until it passes	S. Khan
Coordinators ignore suggestions	M	M	Co-design review screen; show reasons	M. Pillai
Claims API too slow for 5-minute target	L	M	Load test early; 10-minute batch fallback	S. Khan

✦ 05 / PEOPLE

Roles, decisions and timeline

Name people, not teams. A RACI with “the business” in it is a RACI nobody owns.

Decision RACI (R = responsible, A = accountable, C = consulted, I = informed)

DECISION	SPONSOR	PRODUCT OWNER	ENG LEAD	RISK / LEGAL
Scope changes	A	R	C	I
Data access approval	I	C	R	A

DECISION	SPONSOR	PRODUCT OWNER	ENG LEAD	RISK / LEGAL
Model / provider choice	I	C	A	C
Go / no-go for launch	A	R	C	C
Stop or pivot	A	R	C	I

Named roles

01

Sponsor, product owner, engineering lead, data owner, risk / legal reviewer, end-user champion.

WORKED EXAMPLE

- **Sponsor:** COO — V. Iyer · **Product owner:** Head of Claims Operations — R. Menon
- **Engineering lead:** Lead AI Engineer — S. Khan · **Data owner:** DPO — A. Das
- **End-user champion:** Senior Claims Coordinator — M. Pillai

Milestones and phase gates

02

e.g. “Discovery done [date] → prototype on golden set [date] → pilot with [n] users [date] → launch decision [date].”

WORKED EXAMPLE

- Baseline signed off 30 Oct 2026 → golden-set prototype 4 Dec 2026
- Pilot with 6 coordinators 11 Jan 2027 → launch decision 5 Feb 2027

Budget and resourcing

03

Build budget, monthly run budget ceiling, and the people committed (with % of their time).

WORKED EXAMPLE

Build ₹45 lakh; run ceiling ₹1.2 lakh/month. Two engineers full time; R. Menon 30%, M. Pillai 20%.

✦ 06 / SIGN-OFF

Review and sign-off

Tick each item before signing. An unchecked box is a conversation to have now, not later.

- ☐ The problem is stated without mentioning a technology

- [ ] A primary metric, baseline and target are written down
- [ ] Out-of-scope items are explicit
- [ ] Every data source has a named owner and an access route
- [ ] Prohibited behaviors and human-review points are agreed
- [ ] A kill criterion and phase gates are dated
- [ ] Every RACI decision has exactly one accountable person

Decision 01

Approved / approved with conditions / not approved — and the conditions.

WORKED EXAMPLE

Approved with conditions: DPO signs off health-data handling before the prototype uses live claims.

Sponsor — name, signature, date 02

WORKED EXAMPLE

V. Iyer, COO · approved · 12 Oct 2026

Product owner — name, signature, date 03

WORKED EXAMPLE

R. Menon, Head of Claims Operations · approved · 12 Oct 2026

Engineering lead — name, signature, date 04

WORKED EXAMPLE

S. Khan, Lead AI Engineer · approved · 13 Oct 2026

✦ GET IN TOUCH

Want us to run this with your team?

We facilitate charter workshops that leave you with a signed, scoped AI project and a realistic first milestone — usually in a single session.

- ✦ Kickoff workshop facilitation
- ✦ Baseline and metric design
- ✦ Data access and risk review

BOOK AN AI STRATEGY SESSION →

EMAIL
hello@splendenslabs.com

WEB
splendenslabs.com

OFFICE
Sona Incubation Foundation, Salem