

TEMPLATE 01 · FOR FOUNDERS BUILDING AI PRODUCTS

The AI MVP brief: one page that decides what you build first

This is the document I ask every founder for before the first call. Fourteen fields, one page, plain sentences. If you can fill it in, a good team can estimate and start your AI MVP in days. If you can't fill a field, that field is your first piece of work — and it is far cheaper to discover that now than in week five of a build.

45 minutes

14 fields + worked example

Updated October 2026

How to use it

1. Write in plain sentences, not bullet fragments. "Users want AI" is not an answer; "Sales reps lose 40 minutes a day re-typing call notes into HubSpot" is.
2. Use a number wherever you can: minutes, money, volume, error rate. Guesses are fine — mark them with (est.).
3. Keep it to one page. If a field needs more than 3–4 sentences, you are describing the roadmap, not the MVP.
4. Share it with whoever builds the product. A good partner will push back on the cut list and the AI role — that is the point.

Part A — Product

Product in one sentence

For [primary user] who [situation], [product] [does the core job], unlike [current alternative].

Problem and evidence

What is painful today, for whom, how often, and what does it cost? What have you seen or heard that proves it?

Primary user (one) and other people involved

Name one person you design for first. List others who touch the product (buyer, admin, reviewer) but do not design for them yet.

Core job

When [situation], I want to [motivation], so I can [outcome]. Exactly one job for the MVP.

Current alternative and cost of the status quo

What do people do today instead? Why is it not good enough? What does switching cost them?

Part B — The AI

What the model does

The specific transformation: input → output. Be concrete about format.

What the model does NOT do

Draw the line explicitly. What stays deterministic code or a human decision?

Human in the loop and acceptable failure

Where does a person review, edit or approve? What happens when the AI is wrong — and how wrong is acceptable?

Data sources and access

Where does the input come from? Format, volume, permissions, personal data, retention.

Part C — Success and scope

Success metrics

One north-star for the MVP plus: activation, quality (eval), cost per task, and one guardrail that must not get worse.

Scope: must-haves (max 5) and the cut list

Five must-haves for v1. Then an explicit list of what is NOT in v1 — this list is more important than the first one.

Riskiest assumption and how to test it in week 1

What, if false, kills the product? How can you test it before building most of it?

Constraints

Platforms, integrations, compliance (GDPR, HIPAA, SOC 2), languages, accessibility, existing tech you must reuse.

Budget, timeline and decision-maker

Budget range, hard deadline and why it exists, who signs off, and how much of your time is available per week.

Worked example: CallNote

CallNote is a fictional product used to show what good answers look like. Numbers are illustrative.

Product in one sentence

For B2B sales reps who run 5–8 discovery calls a day, CallNote turns each call recording into a ready-to-approve CRM update in under a minute, unlike manual note-taking or generic transcription tools that still leave the CRM empty.

Problem and evidence

Reps spend 30–45 minutes a day writing call notes and updating deal fields. In 14 interviews, 11 reps admitted they skip updates on busy days; sales managers then forecast on stale data. One design-partner team (12 reps) measured 6.5 hours per rep per week on CRM admin.

Primary user (one) and other people involved

Primary: an account executive at a 10–50 person B2B SaaS sales team, uses HubSpot, on calls via Zoom or Google Meet. Also involved: the sales manager (buyer, reads the pipeline), RevOps admin (installs the integration).

Core job

When I finish a discovery call, I want the CRM to already reflect what was said, so I can move to the next call without admin work and my manager trusts the pipeline.

Current alternative and cost of the status quo

Manual notes in a doc, then copy into HubSpot; or Gong/Otter transcripts that nobody reads. Transcription exists — the gap is structured, field-level CRM updates the rep trusts enough to approve in one click.

What the model does

Input: call transcript + existing deal record. Output: a JSON update with 6 fields (next step, close date change, budget mentioned, decision makers, objections, competitor mentioned), each with the transcript quote it came from, plus a 5-line summary.

What the model does NOT do

Never writes to the CRM without the rep's approval. Does not change deal stage or amount (rep-only). Does not send emails to prospects. Does not score or rank reps. Pipeline math stays in code, not in the model.

Human in the loop and acceptable failure

Rep reviews a diff of proposed field changes and approves, edits or rejects each one. A wrong suggestion costs 5 seconds to reject; a wrong silent write would cost trust in the whole product — so there are no silent writes. Target: 80% of suggested fields accepted without edits.

Data sources and access

Zoom and Google Meet recordings via their APIs (OAuth per rep); HubSpot deals via private app token. ~6 calls/rep/day, 30–45 min each. Contains prospect names and emails (personal data) → EU data region, 30-day transcript retention, provider with zero data retention.

Success metrics

North-star: % of calls that end with an approved CRM update within 1 hour (target 60% by week 4 of pilot). Activation: rep approves first update on day 1. Quality: ≥ 80% of fields accepted unedited on a 50-call golden set. Cost: ≤ \$0.15 per call. Guardrail: zero writes without approval.

Scope: must-haves (max 5) and the cut list

Must: Zoom + Meet import; HubSpot deal matching; 6-field suggestion with source quotes; approve/edit/reject UI; Slack ping when ready. Cut: Salesforce, Teams, mobile app, manager dashboards, coaching/scoring, multi-language, custom fields, email drafting. All are "later" — none blocks learning whether reps trust the updates.

Riskiest assumption and how to test it in week 1

Riskiest: reps will trust AI-written CRM fields enough to approve them without re-listening. Week-1 test: run 30 real transcripts through a prompt in a notebook, show the diffs to 5 reps in a Figma prototype, measure accept-without-edit rate. Below 60% → rethink the output before building integrations.

Constraints

Web app (desktop-first), Chrome. HubSpot only for v1. GDPR: EU hosting, DPA with model provider. English only. Team already uses Next.js and Postgres — reuse.

Budget, timeline and decision-maker

Budget \$25–40k for MVP. Deadline: pilot with 3 design-partner teams by 15 January (their Q1 planning). Decision: CEO (me), weekly 60-min review + async in Slack within 24h. Success at the pilot unlocks the seed round conversation.

Want this done for you?

Stepikin Studio designs and builds AI products end to end and ships MVPs in weeks. Get an estimate: stepikin.com/estimate · Book a call: vizi.club/book/aleksey_stepikin