Validated Problem Statement Template

BY JETTHOUGHTS TEAM - MAY 18, 2026 - 8 MIN READ

Template companion to the Mom Test synthesis section. Print after running 10 interviews. Fill in 30 minutes. Show to 2 advisors before building anything.

The synthesis sheet that turns 10 Mom Test transcripts into a build, pivot, or kill decision.

> ***The Validated Problem Statement (one-page, 5 sections)***
>
> 1. ***Who has the problem*** - persona, vertical, strong-signal count
> 2. ***What it costs them*** - hours per week, dollars per month, one sanity quote
> 3. ***What they've tried*** - 3 workarounds + why each failed
> 4. ***Why now*** - trigger event in last 12 months + market shift
> 5. ***How big is the pain*** - average + strong-signal score → BUILD / PIVOT / KILL
>
> Fill in order Mon morning. Send to your 2 readers (one advisor, one peer) by Mon EOD.

Why this exists

Send forty-odd cold DMs to strangers venting about their CRM and twelve answer - two yes, ten honest paragraph-long no's - and you have something that feels like validation. Left as transcripts in a folder labeled `notes`, they lie to you. The ten no's blur together in memory and contradict each other on the page; half of them turn out not to be the persona you are building for. Sit down to name the persona, tally the strong signals, and write the why-now, and the 90-minute synthesis can hand you a *pivot to a different ICP* where memory was telling you *build*. That is a weekend you would have spent prompting Lovable, spent instead on five sharper interviews. This template is the page you fill in to force it.

How to use this

Block 90 minutes on a single morning - about 30 of them go to the writing itself, the rest to re-reading transcripts and scores. Print the template (or copy the markdown version below into a Notion doc). Bring all 10 interview transcripts, your handwritten Q4 scores, your emotional-flag counts.

The order matters. Score first, count second, write the page third. Write the page before scoring and you write the page you wished the calls had returned, not the page the transcripts actually support. The friction of writing the score before the prose is what stops you from talking yourself into the build.

Take the filled page to 2 readers within 48 hours. One advisor (a founder one step ahead of you, or a mentor from a founder community). One peer (another founder still pre-launch). Ask each: “Would you argue with this?” If both nod, you’re done with Module 2. If either picks a fight on the persona, the cost, or the why-now, you have your next 5 interviews to run.

Total time budget: 30 minutes to write, 48 hours to circulate, 1 hour to incorporate the 2 advisor reads. Hard cap at 3 hours total. Beyond that, you’re polishing instead of validating.

The template - copy and paste

Use the markdown block below directly in Notion, or copy it into any doc tool - then score it against the Mom Test Synthesis rubric.

```
================ VALIDATED PROBLEM STATEMENT ================

Founder name: ________________________  Date: ____________

Interview sample: 10 calls, between ___/___ and ___/___.

-------------------------------------------------------------
1. WHO HAS THE PROBLEM
-------------------------------------------------------------

Persona (named, specific):
____________________________________________________________

Industry / vertical (one, not many):
____________________________________________________________

Of 10 interviewees, this many confirmed the problem:  ___ /10

Quote from a strong-signal call (one sentence, verbatim):
"___________________________________________________________"

-------------------------------------------------------------
2. WHAT IT COSTS THEM
-------------------------------------------------------------

Time per week:          _____ hours
Money per month:        $_______
Sanity cost (one concrete moment from a transcript):
"___________________________________________________________"

Cost is consistent across the 10 calls (Y/N): _____

-------------------------------------------------------------
3. WHAT THEY'VE TRIED
-------------------------------------------------------------

Workaround 1: ______________________  Why it failed: _______
Workaround 2: ______________________  Why it failed: _______
Workaround 3: ______________________  Why it failed: _______

Current spend on workarounds: $_______ /month (average)

The workaround is the budget your product would replace.
Name the closest competitor: ______________________________

-------------------------------------------------------------
4. WHY NOW
-------------------------------------------------------------

Trigger event (what changed in last 12 months):
____________________________________________________________

Market shift (regulation, AI cost, platform, behavior):
____________________________________________________________
```

```
If this section is empty, the why-now is missing.
A missing why-now usually means the idea is good and the
timing is wrong - park it, run a smoke-test landing page,
revisit in 6 months.

------------------------------------------------------------
5. HOW BIG IS THE PAIN
------------------------------------------------------------

Average pain score across 10 calls (Q4 + emotional flags):
______ /10

Strong signals (score 7+ with 3+ emotional flags): ___ /10

Decision based on strong-signal count:
 [ ] 7 or more  → BUILD. Test the shape in Lesson 2.6, then Module 3.
 [ ] 4 to 6     → PIVOT. Run 5 sharper interviews.
 [ ] Fewer than 4 → KILL. Find a different problem.

------------------------------------------------------------
ADVISOR SIGN-OFF
------------------------------------------------------------

Advisor 1 name + role: ___________________________________
Advisor 1 says (one sentence): ___________________________
Date: ______________

Advisor 2 name + role: ___________________________________
Advisor 2 says (one sentence): ___________________________
Date: ______________

============================================================
```

What good looks like vs what bad looks like

Three worked examples for the sections founders most often get wrong - 1, 2, and 4. Sections 3 and 5 are direct transcript copying; the same specific-beats-vague rule applies.

Section 1 - Who has the problem

> Bad: *"Founders and small business owners who need help with productivity."*

> Good: *"Pre-seed B2B SaaS founders running their own bookkeeping reconciliation between Stripe, QuickBooks, and a Google Sheet. 8 of 10 interviewees confirmed (industry: vertical SaaS, May 2026 sample). Quote: 'Last Tuesday at 9pm I spent 40 minutes copying Stripe payouts into my QuickBooks ledger. I called my CFO. She did it in 90 seconds. I felt stupid.'"*

The good version names the persona by stage, industry, and the specific workflow. It puts a number on the strong-signal count and includes one verbatim quote. A peer can argue with it: “Are you sure pre-seed is the right segment? Wouldn’t seed-to-Series-A be the buyer?” That argument is exactly the value. The bad version is too broad to argue with, which is why advisors nod and the post-launch failure mode is *crickets*.

Section 2 - What it costs them

> Bad: *“It costs them time and money. It’s a significant pain point.”*

> Good: *“6 hours per week. $800 per month in CFO contractor time. One founder I spoke with paid $1,500 for a SurveyMonkey panel after the bookkeeping pain spiked - the panel returned a stack of useless responses and she did the work herself anyway. Cost was consistent across 8 of 10 calls; 2 were running their own pre-revenue and had zero contractor spend (but 12 hours of personal time per week).”*

The good version uses real numbers from transcripts. The $1,500 panel anecdote is from a specific person. The 6 hours and $800 are averages with the sample’s variance noted. The bad version is unfilled white space dressed up as prose.

Section 4 - Why now

> Bad: *“AI is changing everything. The market is ready.”*

> Good: *“AI inference cost for the document-classification step fell from $0.04 to $0.001 per call between 2024 and 2026 - the unit economics flip at $9/month per seat. Stripe’s automated tax product (live since 2021) signals SMB finance is being deconstructed feature by feature, but bookkeeping reconciliation at pre-seed founder budgets is still manual.”*

The good version names the specific cost number, the specific competitor’s specific product, and the specific gap. It cites a competitor signal that *supports* the timing rather than refuting it. The bad version is filler that means nothing and helps no one.

What to do after

The decision the filled page makes for you:

- **BUILD** if: 7+ strong-signal calls (score 7+ with 3+ emotional flags), a named workaround the customer is already paying for, and a named why-now from the last 12 months. Test the shape in Lesson 2.6 · Build a Clickable Prototype, then move to the Module 3 Product Brief.
- **PIVOT** if: 4-6 strong signals, OR the pain is real but belongs to a different persona than you targeted (e.g., you interviewed founders but the pain lives with their CFOs). Run 5 sharper interviews with the corrected ICP, then refill this page.
- **KILL** if: fewer than 4 strong signals OR no workaround surfaced in the 10 interviews. Find a different problem and write the 200-word post-mortem below.

Then walk the page through these four moves:

- **Get 2 advisor signatures within 48 hours.** Email the page as a PDF. Ask: "Would you argue with this problem statement?" One sentence response is enough. If both say no, you've passed Module 2's checkpoint and you move to the Module 3 Product Brief next.
- **If you landed in the BUILD lane (7+ strong signals), run the 3 pre-orders test before writing code.** Email your 5 strongest-signal interviewees. Ask each for a $500 deposit, a signed letter of intent, or a paid waitlist slot. Three yes-and-paid out of five = build. Zero yes = the 7+ scores were politer than you thought, slide back to pivot.
- **If you landed in the PIVOT lane (4-6), pick the cleanest segment and run 5 sharper interviews using** the Mom Test interview script **again.** Don't rerun the same 10 - they've already given you their honest answer. New segment, new interviews, one week.
- **If you landed in the KILL lane (fewer than 4), write a 200-word post-mortem to your future self.** What ICP did you pick wrong? What why-now did you assume that wasn't true? What workaround did you not learn about until interview 7? The post-mortem is the most valuable artifact from a kill round - it stops you from picking the same wrong target again next quarter. The stop-AI-obsession validation post has the long version of the discipline.

Skip the synthesis page and start prompting straight away and you're the founder who, six weeks later, posts a salvage or rebuild question about a working MVP nobody wants. The page is cheap. The build round you skip because of the page is the expensive thing you avoid.

Built by JetThoughts as part of the From Idea to First Paying Customer curriculum.