

Session 3 Working Guide

Defend the Call: tie your decision to a number, build it, and set up the loop that learns after launch

What today is for

A decision that doesn't move a number is a hobby. Today you'll take one decided bet, tie it to the number your boss already watches, see it built and tracked, and write the one page that defends it.

You'll leave with:

- a defended, ROI-tied product direction
- a weekly check that tells you whether it's working
- a set of prompts you can reuse on every call

Time (ET)	What we do
12:00	Open: where Launchpad landed
12:05	Build your number (hands-on)
12:25	Live build: prototype → tracking → dashboard → weekly check
12:43	Hands-on: your exec one-pager
1:03	Capstone: pitch it and get pushback
1:25	Close: set up your loop this week

Have these open

1. Your AI tool (Claude, ChatGPT, Gemini, or whatever your company uses).
2. Your Session 2 stress-test output, if you have one.
3. Your strategy doc, vision statement or principles. **Or** the Launchpad Session 3 pack.

Pick your lane:

- **A:** You have a live bet and a strategy doc.
- **B:** You have a bet, but only a vision statement or principles. That counts.
- **C:** Use Launchpad. Play the PM defending the call to the Launchpad exec team.

One rule for every prompt today: your AI tool will invent customer quotes and numbers if you let it. Every prompt below tells it not to. When something sounds too good, ask “where did that come from?”

Part 1: Build your number

The four questions

1. **If this works, what will people do differently?** Describe a behaviour you could see, not a feeling.
2. **How would you count that?** This is your *early number*. It moves in weeks.
3. **What number does your boss or client already report?** This is the number your decision gets judged on.
Examples: PMF, retention, expansion, revenue, pricing or GTM.
4. **What result would make you stop?** Write it down now, before anything ships.

No data yet? Write your best guess and how you’d start counting. A guess you’ve written down beats a number you don’t have.

Launchpad, answered:

Question	Launchpad
People do differently	When an admin leaves, someone new takes over the flows
Early number	Flagged accounts where a new owner edits a flow within 30 days
Boss number	Logo churn, 14% → under 10%
Stop if	Fewer than 1 in 4 new admins edit a flow within 30 days

Prompt 1: Build your number

I’m going to describe a product decision I’ve made. Help me tie it to a number, by interviewing me. Ask me these four questions **one at a time**, and wait for my answer before asking the next:

1. If this works, what will people do differently? Push me until the answer is a behaviour you could see, not a feeling.
2. How would I count that? Help me pick one early number that would move within weeks.
3. What number does my boss or client already report, that this should move? (For example PMF, retention, expansion, revenue, pricing or GTM.) Ask me how the early number connects to it.
4. What result would make me stop? Help me write it as a number and a time window.

Don’t answer for me. If I don’t have data, ask for my best guess and how I’d start counting. At the end, give me a four-line summary.

Here’s my decision: [describe your bet in two or three sentences]

Prompt 2: Cost and return, in your approver's math

Here's my decision and my number: [paste your four-line summary from Prompt 1].

Here's what I know about cost and value: [build effort, team, price, ARR at risk, deal sizes, whatever you have].

My approver is [role]. They usually judge decisions by [first-year revenue / logos / payback period / cost per customer / "I don't know"].

Work out the return three ways: in my approver's math, over a longer horizon, and in the number from question 3. Show the arithmetic in plain sentences.

- Label every assumption as an assumption.
- Tell me honestly if it doesn't clear in my approver's math.
- Tell me what would have to be true for it to clear.
- Don't invent numbers. If you need one I haven't given you, ask me for it.

Why three ways? The same bet can be underwater in one kind of math and fine in another. Launchpad doesn't pay back in first-year ARR. It does on three-year value if we keep just over 1 in 4 accounts, and every saved account is a logo the board counts. Say which math you're using. An approver who catches you switching to friendlier math stops trusting the rest.

Part 2: The build

You'll watch this live on Launchpad. You don't need to build it today. These prompts let you do it on your own bet this week, with whatever tools your team uses.

The chain: decided bet → spec → prototype → code your engineers can review → tracking → dashboard → an alert on your stop line → a weekly check that says what to do next.

Step	AI does	You own
Spec and tickets	Drafts them from the bet	The stop rule, and what you're not building
Prototype	Builds the screen	What the user needs to see
Tracking	Wires the events from a plain list	Which steps prove it works
Dashboard	Builds the funnel and the split	The split that tests your theory
Weekly check	Reads the numbers and suggests a move	Whether to keep going, change course or stop

Prompt 3: Spec and tickets

Here's a product bet we've decided on, with the stress-test results: [paste your bet and your Session 2 output, or your four-line summary].

Write a one-page spec an engineer could build from. Include:

- the problem in two sentences
- who it's for
- what we're building
- what we're not building
- acceptance criteria written as things someone could check
- the events we need to track (name, when it fires, why we need it)
- the rollout plan
- the stop rule

Then break it into 3 to 5 tickets, each with acceptance criteria. Keep it plain. If something is missing from what I gave you, list it as an open question. Don't fill it in.

Prompt 4: What to track

Here's my decision and the behaviour I expect to change: [paste].

List the 3 to 5 steps a user goes through, from "this applies to them" to "they did the thing that proves it works." For each step, give:

- an event name in lower_snake_case
- when exactly it fires
- why we need it
- the properties to attach to it (include anything I might want to split by later)

Then tell me what a drop between each pair of steps would mean, and what I'd do about it.

The Launchpad events:

1. `admin_changed`
2. `ownership_prompt_shown`
3. `ownership_accepted`
4. `flow_edited_by_new_owner`

The gaps between them tell you what to fix:

Where it drops	What it means
Flagged, but not shown	A bug
Shown, but not accepted	People don't see it as their job. No feature fixes that; a person has to ask
Accepted, but no edit	They said yes and still didn't do the work. Maybe the flows really are hard to edit

Prompt 5: The prototype

This works in Lovable, Claude, v0, Bolt or whatever your team uses.

Add a feature called “[name]” to [the page]. Keep the existing design.
 [Describe what the user sees, what each button does, and what happens after.]
 Only show it when the feature flag `[flag-name]` is on.
 Track these events with [your analytics tool], with properties [list them]:

- `[event_1]` when [moment]
- `[event_2]` when [moment]
- (and so on)

Don't change anything else.

Test on a copy first. If the build goes wrong, you still have a working version.

Prompt 6: The dashboard

If your analytics tool connects to your AI tool (for example PostHog with Claude):

In my [tool] project “[name]”, build a funnel with these steps in order: [events]. Count unique users, with a [30]-day conversion window, over the last [60] days. Make a second copy broken down by [the property that tests your theory]. Put both on a new dashboard called “[name] weekly”. Then tell me the conversion at each step and overall.

If it doesn't: export the events as a CSV and use the same words with the file attached. Any AI tool can build the funnel table from a CSV.

Set an alert on your stop line. Most analytics tools can email you when a number drops below a line you set. Launchpad's alert fires weekly if the overall rate drops under 25%.

Prompt 7: The weekly check

This is the prompt that makes it a loop. It judges the numbers against a rule you wrote before launch, not against how you feel that week.

You're checking whether "[feature name]" is working. Use [the dashboard name, or the CSV I've attached].

Our stop rule, set before launch: [your stop rule].

Reply in this order, in plain language:

1. **Verdict:** Keep going, Change course, or Stop. One line.
2. **The number:** each step's count, and the overall rate against our stop line. Don't round in our favour.
3. **The biggest drop**, and what it means. [Paste your "what the gaps mean" list.]
4. **Any split that changes the story.** Say how many accounts each side is based on.
5. **One next move**, and what we'd need to see next week to believe it's working.

Only use the numbers in the data. Don't invent reasons customers did things. If the sample is too small to trust, say so.

What it said about Launchpad:

- **Verdict:** Change course.
- **The number:** 49 flagged → 42 shown → 22 accepted → 10 edited. That's 20.4%, under the 1 in 4 line.
- **The split:** 42% when the CSM called (19 accounts) vs 7% when they didn't (30 accounts).
- **Next move:** make the CSM call part of every handoff.

The loop didn't say "kill it." It said "change it."

Prompt 8: Make it run every week

In Claude, with your analytics tool connected, as a scheduled task:

Every Monday at 9am, run the weekly check below on [dashboard name] and send me the result. [Paste Prompt 7.]

Anywhere else: set a Monday calendar reminder, export the numbers, and paste them into Prompt 7. It's less elegant, but it still works.

Part 3: The exec one-pager

This is the page that ties everything together: your signal from Session 1, your stress test from Session 2, and your number, cost and loop from today. If it doesn't fit on one page, you're not ready.

The seven parts:

1. **The decision:** one sentence.
2. **The number it moves:** the boss number and your early number, with today's value if you have one.
3. **Cost and return:** in your approver's math first, then the others.
4. **What we're not doing, and why:** the trade-off, named.
5. **The receipts:** each claim, where it came from, and its verdict.
6. **When we'd stop:** your stop rule, and where you are against it.
7. **The ask:** exactly what you need, from whom, by when.

Prompt 9: Draft your one-pager

Write an exec one-pager for my decision. Use only what I give you below. Don't add facts, numbers or customer quotes that aren't in these inputs.

Inputs:

- My signal summary (Session 1): [paste, or "none"]
- My stress-test results (Session 2): [paste]
- My number (Prompt 1): [paste]
- My cost and return (Prompt 2): [paste]
- My latest weekly check or test result, if any: [paste, or "not launched yet"]
- My approver: [role], who cares most about [what]

Format: one page, seven sections, with these exact headings:

1. The decision
2. The number it moves
3. Cost and return
4. What we're not doing, and why
5. The receipts
6. When we'd stop
7. The ask

Write "The receipts" as a table with three columns: Claim, Source, Verdict (Confirms, Contradicts, Flatters or Silent). Every claim in the rest of the page must appear in that table.

Plain language. Short sentences. No jargon. Where something is missing, write [MISSING: what's needed] rather than filling it in.

Prompt 10: Find the weak spot before your approver does

Here's my one-pager: [paste]. My approver is [role].

1. Point to the single weakest line, the one most likely to get me asked "how do you know?"
2. Ask me the three hardest questions this person would ask, in the way they'd actually say them.
3. Check every number: does it add up, and does it match the receipts table?
4. List any claim that has no source.

Don't rewrite the page. Just tell me where it breaks.

Part 4: Capstone

How it runs:

1. You pitch your one-pager in 3 minutes.
2. We push back as your approver, using the stakeholder views from Session 2.
3. You get 2 minutes of feedback.

What to bring: your one-pager (a draft is fine), and the one number you'd most want your approver to remember.

What we're listening for:

- Is the number one your approver already reports?
- Is the math theirs?
- Is the stop rule real, meaning could a result actually stop it?
- Does every claim have a receipt?

Part 5: This week, and the Q&A session

Set up your loop before the Q&A session:

- ☐ Four questions answered for your real bet (Prompt 1)
- ☐ Cost and return in your approver's math (Prompt 2)
- ☐ 3 to 5 events written down (Prompt 4)
- ☐ A dashboard or a CSV you can check weekly (Prompt 6)
- ☐ A stop line with an alert, or a calendar reminder
- ☐ The weekly check scheduled or on your calendar (Prompts 7 and 8)
- ☐ Your one-pager, pitched to one real person (Prompts 9 and 10)

Bring whatever breaks to the Q&A session. That's what it's for.

The system you'll reuse on every call

Session	What it answers	What AI does	What you own
1. Signal	What's really going on?	Reads and sorts everything	Which signal matters
2. Stress test	Does it hold up?	Sorts evidence against your strategy, plays your stakeholders	The trade-off
3. Defend	Is it worth it, and is it working?	Builds, tracks, watches, drafts	The number, the stop rule, and the call