LIVE COACHING COURSE · FREE · COHORT 1

Product Management with Agentic AI

Learn product management by doing the job. In eight weeks you take your own product from an idea to real users, learn how AI really works, and build the AI agents that take the busywork off your plate.

8 weeks, live | Saturdays from 12 Dec, 10:00 to 12:00 Malaysia time | Online, from anywhere | Free, pay it forward | 10 to 15 seats

Every Saturday, 12 December 2026 to 30 January 2027, 10:00 to 12:00 Malaysia time (UTC+8).

AT A GLANCE

What is this course?

WHAT IT IS	A free, live, cohort-based coaching program in product management and agentic AI
WHO IT IS FOR	Students, engineers and project managers moving into product, and PMs in their first 3 to 4 years
LENGTH	Week 0 setup, 8 weekly live sessions, then Demo Day
SCHEDULE	Every Saturday, 12 December 2026 to 30 January 2027, 10:00 to 12:00 Malaysia time (UTC+8)
TIME EACH WEEK	About 6 hours: 2 hours live plus about 4 hours of your own work
FORMAT	Live online with video on; optional meetups in Kuala Lumpur when available
COHORT SIZE	

10 to 15 people, so everyone gets real feedback

PRICE

Free and pro bono. Pay it forward

COACH

Amit Goel, Chief Product Officer at Moving Walls

WHY THIS PROGRAM, WHY NOW

Product management has changed. The fundamentals have not.

The door into product has moved

LinkedIn has stopped taking applications for its Associate Product Manager program and now points early-career candidates to an Associate Product Builder program. Some interview loops now ask candidates to build a prototype with AI in the room, and "LLM" has become the most-mentioned skill in PM job descriptions (Lenny's Newsletter, Nov 2026). A certificate does not show you can do the job. A working product with real users does.

AI made building cheap, not deciding

Customer discovery, strategy, business models, metrics and judgment matter more than before, because anyone can now ship the wrong thing quickly. You learn those fundamentals every week, then how AI changes the way you do them.

Most PMs drown in busywork

You build a team of AI agents that takes over research synthesis, market watch, metrics updates, status reports and feedback triage, so your time goes to strategy, customers and outcomes.

WHAT YOU WILL BUILD

Two things, side by side

You learn by doing the job, every week, with fast feedback. Teaching is short; practice is long.

01 · YOUR PRODUCT

One idea of your own, taken to real users

Three loops: a rough prototype by Week 2, a properly researched and specified v1 by Week 5, and a product with real users by Week 7. In Week 3 you decide whether to continue, pivot or kill your idea, while changing course is still cheap. Software, hardware, or both are welcome.

`Prototype · Week 2` `Gate · Week 3` `v1 · Week 5` `Real users · Week 7`

02 · YOUR AI HARNESS

A team of agents for the routine parts of the job

Each week you build one AI agent that takes over a routine part of a PM's job, and one test (an eval) that proves it does the job well. In Week 6 you connect them into a harness: shared context about your product, connections to your documents, calendar, tracker and analytics, and a rule that nothing goes out without your approval.

`Research synthesis` `Market watch` `Strategy critique` `Metrics brief` `Status updates` `Feedback triage`

WHAT YOU WILL BE ABLE TO DO

By the end of the eight weeks, you will be able to

1. Find a problem worth solving by interviewing customers and separating what they do from what they say.
2. Decide what to build, and what not to build, using explicit criteria, evidence and a clear business model.
3. Build a working product with AI tools, without needing to write code by hand.

4. Prove it works with a metrics tree, real usage data and honest experiments.
5. Explain how AI models work: parameters, training and inference, tokens, context windows, why models hallucinate, and the trade-offs between models on quality, speed and cost.
6. Give an AI the right context and keep it there across sessions, and cut the tokens (and cost) a task uses without losing quality.
7. Build AI agents, connect them into a harness, and test them with evals so they can be trusted with real work.
8. Run the routine parts of a PM's week (research synthesis, market watch, metrics updates, status reports, backlog and feedback triage) through your agents, with you in control.
9. Write and pitch clearly enough that a sceptic changes their mind.
10. Handle every PM interview format in use today, including AI product sense, technical AI questions and live build rounds.

WHAT YOU LEAVE WITH

A working product
with real users, and the evidence behind every decision you made

Your own AI harness
a set of agents for PM work, with an eval report showing how reliable they are

Real AI understanding
of models, context, tokens, agents and evals, which you can explain in an interview

A public case study
of your product, written for hiring managers

A decision scorecard
eight decisions, the predictions you made, and how well you called them

A demo and a pitch
a 3-minute demo video and a Demo Day pitch reviewed by industry leaders

THE AI FUNDAMENTALS TRACK

You will understand how AI actually works

You cannot direct a tool you do not understand. Alongside the product work, you learn how modern AI works from first principles: what a model is, what it can and cannot do, what it costs, and how to make it reliable. Each concept is shown live, then used the same week in the agent you build.

Week 0 **How models work**

What a large language model is; parameters and weights; training vs inference; model families and sizes; why the same question can get different answers. You will be able to explain why a model can sound sure and still be wrong.

Week 1 **Tokens and context windows**

How text becomes tokens; context windows as the model's working memory; why models do not remember between chats; why they hallucinate; prompt basics. You will see a model invent a quote, then build the eval that catches it.

Week 2 **Context engineering**

Instructions, reference files, examples and structured output; keeping context across sessions with context files, memory and

Week 3 **Tools and the agent loop**

Tool use and function calling, web search, and retrieval. The agent loop: plan, act, observe, repeat. What makes something an agent rather

summaries; why very long contexts degrade; what to leave out.

than a chatbot, and how to ground answers in sources.

Week 4 **Choosing models and AI judges**

Comparing models on your own task for quality, speed and cost; reasoning models vs fast models; using one model to judge another, and calibrating that judge against your own grades.

Week 5 **Token costs and optimisation**

Input and output tokens, cost per task and per user, latency. Trimming context, caching, sending simple steps to smaller models, limiting output, and knowing when plain rules or code beat AI.

Week 6 **Harnesses and MCP**

Agent = model + harness. Instruction files and skills, MCP connectors to your tools, permissions and human approval, sub-agents with their own context, golden test sets and error analysis.

Week 7 **Safety and regression evals**

Non-determinism and regression evals, prompt injection, private data and guardrails, why models flatter you and how to counter it, what changes when a model is updated, and when prompting, retrieval or fine-tuning is the right fix.

Week 8 **Explaining AI in interviews**

Talking through trade-offs between quality, cost, speed and risk; describing your harness and its evals; answering the technical AI questions PM interviewers ask; staying current as models change.

No maths. No coding required. One-page plain-language explainers for every concept, to keep.

CURRICULUM

Week by week

Every week has the same shape: one decision about your own product, the product management fundamentals that decision needs, one AI concept from first principles, one new AI agent with its eval, and a submission due before the next session. Each week also names one thinking trap to watch, drawn from the coach's essay *The Gospel of Product Monk*.

WEEK	YOUR PRODUCT	PM CRAFT	AI CONCEPT	YOUR AGENT

WEEK	YOUR PRODUCT	PM CRAFT	AI CONCEPT	YOUR AGENT
1	Loop 1: rough idea	The role, customers	Tokens, context	**Research synthesis**
2	Loop 1: first prototype	Choose, prototype	Context design	**Build kit**
3	Loop 2: the gate	Segments, distribution	Tools, loops	**Market watcher**
4	Loop 2: strategy	Strategy, pricing	Models, judges	**Red team + judge**
5	Loop 2: spec and v1	Specs, metrics	Token costs	**Metrics brief**
6	Loop 3: real users	Execution, stakeholders	Harness, MCP	**Harness assembly**
7	Loop 3: learning loop	Retention, growth	Safety, injection	**Feedback loop**
8	Loop 3: launch decision	GTM, pitch	Explain your AI	**Interview coach**

WEEK 0

Get ready

On your own, before the first session. Which problems do you want to explore?

Do before Week 1

- Set up your AI assistant (Claude recommended) and one building tool from the tools list.
- Write a two-page note to yourself: why you want to be a product manager, what you think a PM does, and what you would do if not PM. You will reread it in Week 8.
- Bring one to three problems you would like to explore. They can be software, hardware or both.
- Keep a one-week time log of your current work or study, split into thinking, people, admin and making. This is your baseline.
- Start a one-page context file: who you are, the problems you are exploring, and who you think has them. Your agents will read it all program long.

AI fundamentals: what a large language model is and how it works: parameters and weights, training vs inference, model families and sizes, and why the same question can get different answers. A short recorded explainer gets everyone to the same starting line.

Read: *The Gospel of Product Monk* (Amit Goel).

WEEK 1

The job, and your first customers

The decision: which problem is worth testing first?

You will learn to

- Explain what a product manager does, and how it differs from a project manager, a product owner and an engineering lead
- Break a problem down from first principles instead of copying what others do
- Run a customer interview that uncovers real behaviour, not polite opinions

Topics: the product model vs the project model; what AI made cheap (building) and what it did not (deciding what to build); first principles, Socratic questioning and eigenquestions; story-based interviewing; Jobs to Be Done; how hardware and software product lifecycles differ.

AI fundamentals: tokens and context windows; why a model does not remember your last chat; why it hallucinates; the basics of a good prompt (role, task, context, examples, output format). You will see an interview transcript counted in tokens, watch a model invent a quote, then build the eval that catches it.

Your product: write your problem hypotheses without AI, rehearse an interview with an AI customer simulator, then interview 5 real people.

Your agent: research synthesiser. A skill that turns your interview notes into themes, each backed by a direct quote. **Its eval:** every quote it uses must exist word for word in your notes.

In the gym: estimation questions. **Thinking trap:** confirmation bias, asking questions that only confirm what you already believe.

Submit: 5 interview notes, a problem statement, the research synthesiser with its eval result, and your first forecast.

Read: *The Mom Test* (Rob Fitzpatrick); *Inspired*, part one (Marty Cagan).

WEEK 2

Decide fast, build fast (Crit 1)

The decision: which single solution will you test, and what result would kill it?

You will learn to

- Turn interview evidence into a map of opportunities and choose one
- Build a clickable prototype with AI in under an hour
- Test it with users without leading them

Topics: the opportunity solution tree; choosing between opportunities; building to learn vs building to earn; prototype fidelity; usability testing; setting kill criteria in advance.

AI fundamentals: context engineering. How to give an AI the right instructions, reference files and examples; how to keep context across sessions with context files, memory and summaries; why very long contexts degrade; and what to leave out.

Your product: build a prototype of your core user flow and test it with 3 people.

Your agent: build kit. A spec template and a design rules file that any AI builder reads before it starts, so every prototype begins from your context. **Its eval:** your prototype passes a short usability checklist.

In the gym: product sense ("improve this product for this user"). **Thinking trap:** the IKEA effect, falling in love with what you built.

Crit 1: 3 minutes to show your problem, your prototype and what testing taught you.

Submit: prototype link, 3 test notes, a one-page decision memo with your kill criteria.

Read: *Continuous Discovery Habits*, chapters on opportunity mapping (Teresa Torres); "Build to Learn vs Build to Earn" (SVPG).

WEEK 3

Market, segment and distribution

The decision: continue, pivot or kill, and which customers first?

You will learn to

- Make a continue, pivot or kill call against criteria you set before you were attached
- Segment a market and pick a beachhead
- Size a market from first principles rather than from a top-down guess

Topics: segmentation; ideal customer profile (ICP) and beachhead; buyer vs user vs payer; B2B vs B2C buying behaviour; the technology adoption lifecycle and crossing the chasm; market sizing, and why "1% of a $100 billion market" is a trap; distribution channels: where your first customers already are, and how you will reach them.

AI fundamentals: tools and function calling, web search, and retrieval (bringing the right documents into context instead of retraining a model). The agent loop: plan, act, observe, repeat. Grounding answers in sources so you can check them.

Your product: hold your gate review, then interview 10 more people in your chosen segment.

Your agent: market watcher. Your first true agent: it runs on a schedule, searches for competitor and market news, and writes a weekly digest with sources. You learn the anatomy of an agent: model, instructions, tools, memory, loop and schedule. **Its eval:** the share of its claims that carry a correct source.

In the gym: market sizing. **Thinking trap:** anchoring, letting the first number you hear set your thinking.

Submit: gate decision, opportunity brief, evidence log, and the market watcher's first digest with its citation score.

Read: *Crossing the Chasm*, chapters 1 to 3 (Geoffrey Moore); *Competing Against Luck*, chapter 1 (Clayton Christensen).

WEEK 4

Strategy and the business

The decision: where will you win, and how will you make money?

You will learn to

- Write a positioning statement and a short, clear strategy
- Identify what could give your product lasting advantage
- Build simple unit economics, including the cost of any AI or hardware

Topics: vision and value proposition; competition, including the option to do nothing; 7 Powers and moats in the AI era (data, distribution, workflow); business and revenue models; unit economics, customer acquisition cost (CAC) and lifetime value (LTV), AI running costs and a hardware bill of materials; pricing.

AI fundamentals: choosing a model. Comparing models on your own task for quality, speed and cost; reasoning models vs fast models; using one model to judge another, and calibrating that judge against your own grades.

Your product: write a two-page strategy memo in plain prose, without AI.

Your agent: red team and judge. A critique loop. One agent argues the strongest case against your memo; another scores it against your own rubric. You grade 10 sample memos yourself and adjust the judge until it agrees with you. **Its eval:** how often the judge's score matches yours.

In the gym: strategy questions. **Thinking trap:** optimism bias and the planning fallacy.

Submit: strategy memo, unit economics sheet, red-team log, judge agreement score.

Read: *7 Powers*, introduction (Hamilton Helmer); *Good Strategy Bad Strategy*, chapters 1 to 2 (Richard Rumelt).

WEEK 5

Specs, metrics and v1 (Crit 2)

The decision: what goes into v1, what stays out, and how will you know it worked?

You will learn to

- Choose a north star metric with the input and guardrail metrics around it
- Write a PRFAQ and use a prototype as the specification
- Talk with engineers about front end, back end, APIs and data with confidence

Topics: outcomes over output; north star, input and guardrail metrics; the AARRR funnel (acquisition, activation, retention, referral, revenue); the GAME model (goals, actions, metrics, evaluation); prioritisation frameworks such as RICE, and build vs buy; the PRFAQ; enough technical foundations to work well with engineers; one-way and two-way door decisions.

AI fundamentals: token economics and optimisation. Input and output tokens, cost per task and per user, and latency. How to cut cost without losing quality: trim context, reuse it with caching, send simple steps to smaller models, limit output length, and use plain rules or code where AI is not needed. You will add the AI cost line to your unit economics.

Your product: build v1 with an AI builder or coding agent that reads your spec, and add analytics.

Your agent: metrics brief. Connected to your analytics, it writes your weekly update: what moved and what to check. **Its eval:** every number matches the dashboard, and every claimed cause is labelled as a hypothesis.

In the gym: metric diagnosis ("sign-ups fell 10%; why?"). **Thinking trap:** treating easily reversed decisions as if they were permanent.

Crit 2: 3 minutes to show v1, your metrics tree and your strategy.

Submit: v1 link with analytics, a one-page PRFAQ, a metrics tree, the metrics brief agent.

Read: *Working Backwards*, chapter on the PRFAQ (Colin Bryar and Bill Carr); *Measure What Matters*, part one (John Doerr).

WEEK 6

Execution and stakeholders

The decision: which parts of your week will agents run, and where must you stay in the loop?

You will learn to

- Plan delivery as outcomes, not a list of features
- Assemble your agents into a harness you can trust
- Test agents properly and fix failures in the harness rather than blaming the model

Topics: outcome roadmaps; delivery rhythm; working with designers and engineers; stakeholder management and updates; decision logs; saying no; risk management.

AI fundamentals: how a reliable agent system is put together. Agent = model + harness. Instruction files and skills; MCP connectors; permissions and human approval steps; sub-agents that work with their own context; golden test sets and error analysis.

Your product: plan the next two weeks of v2 and recruit real users.

Your agent: the harness. Bring your agents together with a shared context file (product, customers, decisions), connections to your documents, calendar, task tracker and analytics, and skills for status updates, meeting notes into action items, and backlog triage. Nothing is sent without your approval.
Its eval: a golden set of 20 real tasks with expected results; you run it, read every failure, and fix the harness.

In the gym: an AI product sense question with a 20-minute live build, the format some companies now use in interviews. **Thinking trap:** automation bias, trusting output because a machine produced it.

Submit: harness v1, eval report, and a second one-week time log.

Read: "The Anatomy of an Agent Harness" (LangChain); "Building Effective Agents" (Anthropic); *Beyond the Prompt* (Amit Goel, Colliding Neurons).

WEEK 7

Real users and growth

The decision: what are real users telling you, and what will you change?

You will learn to

- Run an experiment, and learn when you cannot (small samples, hardware)
- Combine what users say with what they do
- Keep agents safe and honest as they touch real user data

Topics: activation and the "aha" moment; retention cohorts; growth loops; experiments and A/B tests; combining qualitative and quantitative evidence; outcome reviews; privacy and responsible AI; influencing people who disagree with you. Optional module: when AI belongs inside your product, and how to evaluate it there.

AI fundamentals: keeping AI reliable and safe over time. Non-determinism and regression evals; prompt injection; private data and guardrails; why models flatter you and how to counter it; what changes when a model is updated; and when prompting, retrieval or fine-tuning is the right fix.

Your product: get at least 5 real users and ship your v2 changes.

Your agent: feedback loop. It gathers user feedback and usage data, groups it into themes, and proposes ranked experiments. You decide which to run. You also learn the risks: prompt injection through user content, private data, runaway cost and over-automation, and you write your own list of what you will never delegate. **Its eval:** after every change to your harness, re-run every eval; no new failures allowed.

In the gym: behavioural questions using STAR. **Thinking trap:** AI sycophancy; an assistant that always agrees is a mirror, not a critic.

Submit: v2, usage data, the feedback agent, a regression report, your "never delegate" list.

Read: Hamel Husain and Shreya Shankar's free AI Evals Email Course.

WEEK 8

Launch, pitch and your next job

The decision: launch, pivot or kill? Make the call with your data.

You will learn to

- Plan a launch that reaches the right first customers
- Pitch with a clear story where every claim traces to evidence
- Review your own decisions honestly and measure your judgment
- Present your work in every type of PM interview

Topics: go-to-market (beachhead, channels, pricing, launch tiers; certification, inventory and retail for hardware); narrative and the pitch template; decision retrospectives; the job search: portfolio, interview formats, your first 90 days as a PM.

AI fundamentals: explaining AI like a product leader. Talking through trade-offs between quality, cost, speed and risk; describing your harness and its evals; answering the technical AI questions PM interviewers ask about context windows, tokens, cost and evaluation; and staying current as models change.

Your product: write your launch plan and prepare your Demo Day pitch.

Your agent: interview coach and sceptical investor. Two rehearsal agents that run mock PM interviews and fire the hardest investor questions at your pitch. You also compare your final time log with Week 0: how much time moved from admin to strategy and customers?

In the gym: a full mock interview with your review partner. **Thinking traps:** hindsight bias, and sounding smart instead of being clear.

Submit: launch plan, pitch deck, public case study, final harness eval report, decision scorecard, and a reread of your Week 0 note.

DEMO DAY

Show your work

Present to a panel of product leaders

Each participant has 6 minutes: a live product demo, real usage data and your launch decision, then 1 minute showing your harness at work with its eval scores. Two guest critics from the industry give feedback and score your work. The day closes with a career clinic.

HOW IT WORKS

How the program runs

Format	Live online, with video on. Optional in-person meetups in Kuala Lumpur when available
Length	Week 0 setup, eight weekly sessions, then Demo Day
Session time	Every Saturday, 12 December 2026 to 30 January 2027, 10:00 to 12:00 Malaysia time (UTC+8) Approximate until the final time is announced
Live sessions	2 hours each week; Demo Day is about 3 hours
Your own work	About 4 hours a week, so about 6 hours in total
Cohort	A small group of 10 to 15 people
How you work	On your own product, alone, with a different review partner from the cohort each week
Feedback	Written feedback on every weekly submission within 48 hours, plus an optional 30-minute office hour each week
Community	A cohort WhatsApp community for announcements and peer discussion
Cost	Free. Pay it forward

INSIDE EACH 2-HOUR LIVE SESSION

0:00
The gym
A quick drill with no AI allowed

0:15
Teach by example
One concept, through a real document and decision

0:40
Agent lab
This week's AI concept, shown live

1:00
Studio
Work on your decision with the coach rotating

1:35
Decision clinic
One decision; everyone predicts the outcome

1:52
Commit
Your forecast and next actions

In Weeks 2 and 5, part of the session becomes a crit: each participant gets 3 minutes to show their work and receive direct feedback.

TOOLS

Claude is recommended and used in live demos. Use any AI assistant you prefer, such as ChatGPT or Gemini; the methods are the same. Free tiers work for most of the program; a paid plan helps in Weeks 5 to 8. For building without code: Claude artifacts, Lovable, Bolt, v0 or Replit. Tools change quickly, so an updated list is shared at the start of each cohort.

BRAIN FIRST, AI SECOND, EVIDENCE THIRD

You write the first draft of every decision yourself, because people who let AI do their thinking during practice do worse once it is taken away (Bastani et al., PNAS, 2025). Then AI critiques, challenges and extends your draft. Real customers and real data settle the question; AI-simulated customers are for rehearsal only and never count as evidence.

READY FOR THE MARKET

Every week doubles as interview and job preparation

The program is built around what employers now test for: interviews, your first months in a real product job, and building real products of your own.

Product sense interviews

EMPLOYERS LOOK FOR
Clear thinking about users, problems and trade-offs

HOW YOU PREPARE

AI product sense and live build rounds

EMPLOYERS LOOK FOR
Designing a product, then building a prototype with AI in the room

HOW YOU PREPARE

A drill every week, and eight real product decisions you can talk through

Weekly builds, and a timed 20-minute build drill in Week 6

Technical AI questions

EMPLOYERS LOOK FOR

Understanding of models, context windows, tokens, cost, evals and agents

HOW YOU PREPARE

The AI fundamentals track, and agents you built and tested yourself

Metrics and execution rounds

EMPLOYERS LOOK FOR

Choosing metrics, diagnosing changes, planning delivery

HOW YOU PREPARE

Your own metrics tree, metrics agent and weekly metric drills

Behavioural rounds

EMPLOYERS LOOK FOR

Stories of real decisions, conflict and failure

HOW YOU PREPARE

Your decision journal, gate call and Demo Day give you true stories

Your first months as a PM

EMPLOYERS LOOK FOR

Getting up to speed fast and freeing time for strategy

HOW YOU PREPARE

A harness you can point at real work on day one

YOUR COACH

Learn from someone who does the job



Amit Goel

Chief Product Officer, Moving Walls · Kuala Lumpur

Amit is Chief Product Officer at Moving Walls, a global digital out-of-home advertising technology platform preparing for an IPO. He owns product strategy and the roadmap, and runs

an AI research and development team building MCP servers and agent orchestrators: the same kind of agentic systems you will learn to build in this program.

He has spent 24 years in technology: 11 as an engineer and more than 12 in leadership, across six industries and on both sides of the ad auction. He has also built products of his own, Penome, a collaboration tool for product managers, and Nivritee, a free personal finance app with an AI assistant, and he writes long-form analysis in BidStream and Colliding Neurons. He taught an earlier version of this course, *Learning Product Management*, in the same Pay It Forward spirit.

> Code either runs or it does not.

> Does this move the business?

- 2024 to present
  Chief Product Officer, Moving Walls
- 2021 to 2024
  Lead Principal PM and APAC Product Head, The Trade Desk; retail media with major Fortune 500 retailers
- 2017 to 2021
  Director and Head of Product, Ad-Tech, Amagi; built the OTT ad platform from zero to about a billion ad requests a month
- 2015 to 2016
  VP and Head of Product, Knowlarity Communications; early AI-driven voice analytics
- 2012 to 2015
  Founder and CEO, Patterbuzz; raised angel funding and launched two first-in-India products
- 2005 to 2012
  Engineer, Connected TV, NDS (now Synamedia); set-top box software for DirecTV, BSkyB, Canal Plus, Tata Sky and Airtel Digital TV
- 2001 to 2005
  Software Engineer, APT Software and Rebaca Technologies

amitgoel.me LinkedIn X @amitgoel78 BidStream Colliding Neurons

GUEST CRITICS

Feedback from people who hire product managers

Two product leaders from the industry review your work at the crits and score Demo Day. Names will be announced before the program starts.

To be announced

Product leader, guest critic

To be announced

Product leader, guest critic

WHO IT IS FOR

Is this course for you?

This is for you if you are

- A computer science or engineering student who wants to move toward product
- An MBA student aiming for a product role in a tech company
- An engineer or project manager in a tech company who wants to make the switch
- A product manager in your first 3 to 4 years who wants a stronger foundation and real fluency with AI agents

Software, hardware, or a product that combines both: all are welcome. No coding and no prior product experience are needed. You need a laptop, a reliable internet connection, an AI assistant account, and a willingness to talk to real people and show unfinished work.

This is not for you if

- You want a certificate without the work
- You want to learn one specific tool
- You cannot give about 6 hours a week for eight weeks

Free. Pro bono. Pay it forward.

There is no fee now and no catch later. The only ask: if it helps you, help someone else learn. Coach a participant in a future cohort, review at a crit, publish your case study so others can learn from it, or teach what you learned to someone at work.

FREQUENTLY ASKED QUESTIONS

Questions people ask

How much does it cost?

Nothing. The program is free and pro bono, run in a Pay It Forward spirit: if it helps you, help someone else learn.

What time are sessions in my country?

Every Saturday from 12 December 2026 to 30 January 2027, 10:00 to 12:00 Malaysia time (UTC+8). That is Friday evening in the Americas, very early Saturday in Europe and 07:30 to 09:30 on Saturday in India. That is Friday evening in the Americas, very early Saturday in Europe and Saturday morning in India. The time is approximate until it is announced.

When will dates be announced?

Before the cohort starts, by email and in the WhatsApp community, with a calendar file you can add in your own time zone.

How much time does it take?

About 6 hours a week for eight weeks: a 2-hour live session plus about 4 hours of your own work (readings, customer interviews, building and your weekly agent).

Do I need to know how to code?

No. You will build with AI tools that work from plain language. Engineers who want to go further can use coding agents.

Do I have to use Claude?

No. Claude is recommended and live demonstrations use it, but ChatGPT, Gemini or another assistant works too. The methods are the same.

Will I need to pay for any tools?

Free tiers cover most of the program. A paid plan for your AI assistant helps in Weeks 5 to 8, when your agents do more work, but it is optional.

Can I work in a team?

No. Everyone builds their own product, so everyone practises every part of the job. You will still work with a different review partner each week.

Can I bring a product from my job?

Yes, if you change sensitive details and keep your agents away from company systems unless your employer agrees.

What if my idea turns out to be weak?

That is what the Week 3 gate is for. Pivoting or starting again at that point is normal and encouraged.

Can my product be hardware?

Yes. You will prototype the app or a simulated version of the device, and work through hardware economics and launch.

Does my product need to use AI?

No. AI is how you work as a PM, not a feature your product must have. If your product does use AI, an optional module in Week 7 shows you how to evaluate it.

What if I miss a session?

You can miss one and still complete the program: attend at least 7 of the 8 live sessions, submit at least 7 of the 8 weekly submissions, and present at Demo Day. Tell the coach in advance and submit that week's work as usual.

Who will see my phone number on WhatsApp?

The cohort runs as a WhatsApp Community. Official updates go to an announcement group where only the coach and admins can see your number. A separate discussion group is optional; if you join it, members there can see each other's numbers unless you use a WhatsApp username.

Can I still register if Cohort 1 is full?

Yes. Registration never closes. If Cohort 1 is full when you confirm your email, you keep your place in the queue for the next cohort and are offered any Cohort 1 seat that opens up.

Is there a certificate?

The program gives you something more useful to an employer: a working product, your agent harness with its eval report, a public case study, and written feedback from industry critics.