AI 'bugs' become a quality bar to define, and AI can't write your strategy for you

PM Daily Digest

2026-11-25

AI 'bugs' become a quality bar to define, and AI can't write your strategy for you

By PM Daily Digest • November 25, 2026

Teresa Torres and Petra Wille on redefining bugs for AI products, two Mind the Product talks on keeping strategy authorship human, Match Group's six-message north star metric, and fresh PM hiring data.

In AI products, a "bug" is a gap against a quality bar

Petra Wille and Teresa Torres released a podcast episode on how the idea of a bug breaks down in AI products. Torres builds AI interview snapshots and opportunity solution trees with Vistaly. She finds that an agent's output "isn't right or wrong so much as better or worse," and that deciding what counts as "good enough" is the hardest call of all [1]. The work is defining "good" across several dimensions, such as an email's tone, length and fit for the recipient [2].

Their main warning is about overfitting. If you patch the prompt for every complaint ("never use em dashes"), the agent becomes brittle and does worse on cases it hasn't seen [2, 3]. Instead, fix categories of errors: collect enough examples, run error analysis, and push error rates down over time [3]. This also changes what PMs can promise. Sales, support and the rest of the company need a shared idea of "good enough," and they need to stop expecting that every bug gets fixed [4]. The episode builds on last brief's eval-bar framework by adding the customer-facing side of it.

Strategy: keep authorship, and test the assumptions underneath

Two Mind the Product sessions make the same point: AI makes you faster, but it doesn't make your strategic choices for you.

- **Write the first draft yourself.** In a strategy interview, the guest defines strategy as "an evolving theory of winning": a causal account of why customers choose you over alternatives and keep choosing you. Roadmaps, goals and mission statements are stand-ins for strategy, not the real thing [5]. One product leader fed their top 50 priorities to an AI and got back a polished strategy that "was making no choices." Six weeks later it had vanished "like an etch a sketch" [5]. The guest's rules are to keep authorship of the choices and to write the first draft yourself, because LLMs are "phenomenal editors and terrible ghost writers" [5]. Practical steps: start with an hour-long sketch of the implicit strategy, meaning what the organization actually does today. Then compare it with alternatives, and revisit the strategy roughly quarterly [5].
- **Go up the stack.** Martin Eriksson argues that most companies treat AI as a feature or a way to cut costs, when it should prompt them to rethink strategy. His example is Chegg, which launched an AI version of its product without asking whether students would keep paying for answers they could get free. Its revenue has since fallen 75% [6]. Intercom questioned its seat-based model instead and moved to 99¢ per resolution. Eriksson says its AI agent, Fin, passed $100M ARR in its first year [6]. His advice: name the assumptions at each layer, plan for more than one future, and write down your bets and kill criteria. Be willing to move at 70% confidence [6].

Match Group's north star metric is a "spark"

Spencer Rascoff says Match Group now treats people leaving its apps as success [7]. The company used AI to analyze millions of chats and found that conversations get past small talk after six or more back-and-forth messages. It calls that a "spark" and aligned recommendations, lifecycle marketing and business goals around it [7]. The team now runs a roadmap built specifically to drive more sparks [7]. It also limits what AI can do: the apps suggest topics but won't draft messages, and they won't let outside agents swipe or reply for users [7].

Qualitative signal over telemetry

Obsidian's kepano says app telemetry would never have shown that a browser web clipper was worth building. A user's complaint on X prompted it, and two years later the extension has millions of users [8]. Hiten Shah points to this as the case for being "data-informed versus data-driven," where data includes qualitative evidence [9].

Hiring signals

- A recruiter's sample of 710 US PM postings from the last 30 days included only 9 entry-level roles; 64% were senior or above. 36% of titles mention AI,

and the median posted pay for AI-titled roles was $172.5k versus $163.5k for the rest [10].

- Aakash Gupta counts 316 product people at 22 companies that train AI models, and 34 open "model PM" roles. Tencent has 7 of those openings, more than OpenAI and Anthropic combined. 76% of the postings he read in October asked for evals experience [11].

Smaller signal: who owns the customer's voice

A product marketer at a regulated B2B company was denied access to call recordings because sales called them "our private calls" [12]. The most useful reply was to push for a shared customer-evidence database, not for raw call access. It could hold approved transcripts, redacted snippets, tagged objections and win/loss reasons [13].

Sources

1. X post by @ttorres
2. X post by @ttorres
3. X post by @ttorres
4. X post by @ttorres
5. Your roadmap isn't a strategy: how to build one that sticks
6. Welcome to the software renaissance: Martin Eriksson at Mind the Product Chicago 2026
7. Match Group's CEO wants you to delete his apps (with Spencer Rascoff & Maya Shankar) | Summit '26
8. X post by @kepano
9. X post by @hnshah
10. r/ProductManagementJobs post by u/Disastrous-Pen1340
11. substack
12. r/ProductMarketing post by u/pavanidiotic
13. r/ProductMarketing comment by u/CockroachInCrocs