Cagan: AI made feature teams faster, not better — the bottleneck is deciding what to build

PM Daily Digest

2026-11-17

Cagan: AI made feature teams faster, not better — the bottleneck is deciding what to build

By PM Daily Digest • November 17, 2026

Marty Cagan on why AI raises the stakes on the product operating model, plus new hiring data favoring PMs over design and UX research, wide variation in PM-to-engineer ratios, and practical discovery and roadmap advice.

AI sped up feature teams without fixing outcomes

In a conversation with Joe Murray, Marty Cagan argues that companies running a stakeholder-driven feature roadmap used AI to "turbocharge that same process." Building got fast, but they are "going faster than ever, but not seeing improved results" [1]. His diagnosis: the real problem has always been figuring out what to build, and the cost of development hid that [1]. In the product operating model, leaders assign problems with outcomes, and empowered teams find solutions that are valuable, usable, feasible and viable. AI helps those teams with discovery as well as building [1].

Points PMs can use: - **Value and viability are the risks that sink efforts.** Value means solving the problem better than competitors do, so customers actually switch. Viability has to hold up for legal, compliance and sales, not only for the stakeholder who asked [1]. - **Treat synthetic users with care.** AI lets teams prototype and test faster than before. But synthetic-user results are "just as good as the source," and Cagan says there's "a lot of garbage out there" [1]. - **He doesn't believe in the "SaaSpocalypse."** Serious SaaS apps contain thousands of business rules, often undocumented, and he hasn't seen a single predicted AI-built ERP replacement in three years [1]. Instead, vendors are opening their platforms to agents, which he thinks makes them more valuable [1]. Customers weighing a build-it-yourself option often count only development cost and ignore operating cost [1].

Hiring data favors PMs; UX research keeps falling

Lenny Rachitsky's fifth twice-yearly job market report finds that PMs are "winning in the AI era." It also finds that engineering demand has stalled, design openings are at a record low, every role now requires AI skills, and remote roles are at a record low [2]. A follow-up chart shows data science openings more than doubled from their 2023 lows, while user research openings are down 85% from 2022, a record low [3]. Chamath Palihapitiya reads the data as showing that narrow skills are in less demand than "the ability to solve more generalized, abstract problems" [4].

Aakash Gupta shared Levels.fyi submission data from Zuhayeer Musa as a directional proxy, not a headcount. It shows about 30 engineers per PM at Apple, 12 at Google, 6.6 at Microsoft, 3.6 at Walmart, and 2.6 and 1.8 at Amex and Mastercard [5]. Gupta guesses that coordination and compliance work pushes the payments companies toward more PMs. He calls AI "the big variable" in where the ratio goes from here [5].

Discovery: ask about the struggle, not "why"

A PM building an opportunity solution tree for a dev-tools product found that all 12 interviews produced feature requests like "needs bulk export." Their plan was to ask "why" on every sticky note [6]. A reply pushed back: the tree exists to find underserved needs, not to reword requests until they sound like problems. Instead, have users walk you through what they were trying to do, what got in the way and how they worked around it. Keep the tree going as long as the outcome matters, and prune it as you learn [7]. This month, Teresa Torres's group read of *Continuous Discovery Habits* covers Chapter 11, on choosing analytics and tying them back to the outcome [8].

Roadmap access is a trust problem

A PM whose VP wants an internal roadmap that can't be shared or screenshotted got little sympathy on r/ProductManagement. One commenter said a roadmap is a communication tool and asked what the VP actually wants to achieve [9]. Another called it fear winning over trust. Their fix: label any display with the date it was made and how often it gets reviewed [10].

AI output that's almost good enough

After 20+ years of launching software, Hiten Shah says he nearly shipped an AI-made launch video that barely explained the product. He then studied 8 recent launches and wrote a playbook [11]. A B2B cybersecurity PMM described the same risk in messaging: "almost good" AI assets look fine to non-experts, so they get used. Fast production also leads to sprawling specialized messaging, and the long central guide now mostly serves as a source for AI tools rather than something people read [12].

Briefly

- **Muse:** The original team was 10 people, kept small on purpose so communication and decisions stayed tight [13]. Sarantakos says she can now go from idea to built prototype to real discussion "in a 24-hour cycle" [13].
- **Agent coordination:** On a16z, Alex Imas cited benchmarks where one agent beats two agents splitting the same task, while two people outperform one person with double the resources. Today's agent swarms mostly run tasks in parallel rather than coordinating [14].

Sources

1. Product management in the age of AI: Marty Cagan and Joe Murray
2. X post by @lennysan
3. X post by @lennysan
4. X post by @lennysan
5. substack
6. r/ProductMgmt post by u/parker_hughes_85
7. r/ProductMgmt comment by u/DeanOnDelivery
8. X post by @ttorres
9. r/ProductManagement comment by u/Bernhard-Welzel
10. r/ProductManagement comment by u/Devlonir
11. X post by @hnshah
12. r/ProductMarketing post by u/thexyzaffair
13. X post by @lennysan
14. What History Tells Us About the AI Revolution | Tyler Cowen & Alex Imas