Zapier, Ramp and Replit rebuild around agents as the PM work shifts to deciding what to build

PM Daily Digest

2026-11-19

Zapier, Ramp and Replit rebuild around agents as the PM work shifts to deciding what to build

By PM Daily Digest • November 19, 2026

Operators at Zapier, Ramp and Replit describe products and teams rebuilt around agents, and Postman reports PMs prototyping instead of writing PRDs. Also: a context-sharing argument for AI workflows, a free-tier pricing case study, and a viral-sharing test.

Three operators on rebuilding around agents

Zapier. Wade Foster says Zapier MCP is now Zapier's fastest-growing area. Users install it in the agent they already use, and the agent builds their automations in code. He calls Zapier "a code-based product now, even though the users aren't" [1]. His test for an agent-first product: find what agents are still bad at that the models and harnesses won't fix. Then build a tool that an agent can sign up for, buy and use the way a person would [1]. On cost, he tells customers to have the agent build a workflow rather than act as one. Deterministic Zaps run about 71% cheaper in some benchmarks than an agent that reasons from scratch on every run [1]. His view on competing with the AI labs: the best companies quickly drop experiments that new models make pointless [1].

Ramp. Ramp counted 7,000 pull requests from non-engineers in Q2, mostly from its customer-experience team. Two years ago that number was "functionally zero" [2]. The speaker contrasts this with the usual path, where a ticket goes from PM to designer to engineer and ships two months later, if it isn't marked "won't fix" [2]. Ramp says its MCP usage is growing about 10% per week, and it began onboarding customers to its Glass harness this week [2].

Replit. When vibe-coding competitors spent heavily on ads, Amjad Masad chose "unsexy" enterprise work instead. Replit asked customers like American Express what would make them comfortable building mission-critical apps on

it, and that enterprise business has compounded [3]. With agents handling task-level work, he says staff now spend more time with customers and on "what's worth building" [3].

The common thread: agents are becoming customers for PMs to design for, as well as tools that change who on a team ships.

The PM role: prototypes over PRDs

Aakash Gupta passes on Postman's founder saying that PMs there are "building APIs and prototyping in Claude instead of writing PRDs," and that designers ship PRs through Cursor. Gupta calls it a harder job that needs taste and technical instinct [4]. Market data points the same way. Juicebox's data on recruiter outbound searches shows falling demand for software engineers, rising demand for PMs and faster growth for AI engineering. Design hiring is down 20% year over year. Lenny Rachitsky says this matches his own job-market report [5]. Gupta's Levels.fyi analysis covers 5,013 PMs and 47,318 engineers. PMs start $25K behind engineers in pay, catch up around year 8 ($256K vs. $260K) and earn more after year 12 [6].

Context is built through the work

In TBM 452, John Cutler lists ways context-sharing fails even among competent people. Experts can't put tacit judgment into words, the people asking can't say what decision they're trying to make, and clarifying questions come across as tests of competence [7]. He questions the hype around "giving AI context." Instead of making ever more knowledge machine-readable, he suggests designing AI that builds context the way people do, through collaboration, observation and doing the work together [7].

A thread where PMs asked how to get better at prompting landed in a similar place. The top reply calls poor results "a clarity of thought issue," not a prompting problem [8]. Another commenter warns that too much context, or outdated context, can ruin a session [9]. A third keeps a doc tree on how their systems work that the agent reads. The agent researches and records what it finds before it writes any ticket, so each new chat doesn't start from zero [10].

Pricing: high conversion hid one-time demand

A music-library transfer app found that its free limit converted well, but most subscribers cancelled after their one transfer. After a decade of flat revenue, it made basic transfers free and now charges only for auto-sync and playlist tools. Revenue "dropped significantly" [11]. One reply argues that a flat per-transfer price would have kept the one moment users clearly paid for [12]. Another says to pick leading signals now, such as second-month return rates and sync conversions, and set a date to review them [13].

Sharing that actually spreads

In an r/startups thread, one commenter offers a test for viral sharing: can the person who receives the shared item get value in under a minute, and is your app the obvious way to get more? Quizlet study sets pass. Streaks mostly help retention, not acquisition [14]. Another suggests measuring how many recipients open a shared result without signing up. If almost none do, the share "is just noise" [15].

Sources

1. Zapier CEO on Rebuilding Products to Be Agent-First | Wade Foster
2. Ramp's Bet on the 100x Employee
3. From near-collapse to $9 billion: Replit's Amjad Masad on riding the AI boom | Summit 2026
4. substack
5. X post by @lennysan
6. substack
7. TBM 452: Contexting vs. Context
8. r/ProductManagement comment by u/TheOneTrueEris
9. r/ProductManagement comment by u/NeCede_Malis
10. r/ProductManagement comment by u/metajames
11. r/startups post by u/bimusiek
12. r/startups comment by u/gradyandersen
13. r/startups comment by u/SirWeary7070
14. r/startups comment by u/amaan630
15. r/startups comment by u/techtuber101