

AI Product Work: Faster Evidence, Stronger Guardrails

PM Daily Digest

2026-07-24

AI Product Work: Faster Evidence, Stronger Guardrails

By PM Daily Digest • July 24, 2026

AI product work is most valuable when it shortens the path from customer evidence to tested solutions. This brief covers agent guardrails, a support-ticket discovery loop, Hertility’s high-stakes AI workflow, and a nuanced framework for evaluating managers.

Big Ideas

Use AI to accelerate evidence—not just output

The central distinction for AI-enabled product teams is whether they are accelerating old delivery artifacts or accelerating discovery of solutions that satisfy **customer value and company viability**, then testing those solutions with users and stakeholders. [1] That matters because deep customer intimacy still comes from experiencing the product and its problems as customers do—not from treating customers as abstractions. [2]

Apply it: Define AI work around a learning question: *What customer problem or business constraint will this help us validate?* Use AI to synthesize evidence and prototype options, then put those options in front of customers.

Agent guardrails must be architectural

Persistence, creative problem-solving, and shortcut-finding make agents useful—but can also make them unsafe when goals are misaligned or their environment is insufficiently contained. Prompt restrictions are conventions, not hard boundaries. [3]

Apply it: For agents that touch APIs, databases, or third-party services, define permissions, containment, and approval points in the system design. Do not rely

on instructions in a prompt as the primary control.

Tactical Playbook

Turn support noise into a discovery-and-action loop

1. **Aggregate the raw signal.** Summarize each support conversation into the customer question, root cause, and resolution; cluster recurring themes rather than asking PMs to read thousands of tickets. [4]
2. **Map themes to accountable teams.** Attach product, customer, and taxonomy metadata so an issue can be assigned to the relevant PM and product area. [4]
3. **Prioritize a small number of root causes.** Treat each cluster as a potential “bucket of work”; start with a limited set of high-impact issues rather than attempting to fix every ticket category. [4]
4. **Test the smallest viable version.** Ask for the project’s “V0”: the minimum version that reveals whether it will work. Reducing time per initiative enables more experiments and limits investment in ideas that fail to deliver impact. [5]
5. **Keep human decision gates in delivery.** In one enterprise coding-agent workflow, every key phase paused for developer approval, with a separate model reviewing pull requests. [4]

Why it matters: This connects customer evidence to roadmap ownership while preserving judgment over both what gets built and what gets released.

Case Studies & Lessons

Hertility: outcome-led AI in a high-stakes workflow

Hertility began with a concrete outcome: reduce the time to diagnosis in women’s health, where conditions such as endometriosis can take close to **10 years** to diagnose and primary-care triage can require **10+ appointments**. [6] Its products use a Bayesian diagnostic model to give clinicians probability-based assessments with reasoning, and scan automation to assist image analysis and clinical-letter drafting. [7]

The design is explicitly assistive rather than substitutive: clinicians remain in the loop, and patient-facing outputs receive a two-step clinician review. [6] The team also uses labeled data, holdout sets and independent review to mitigate automation bias; clinicians can correct image contours and labels to support improvement over time. [6] Regulation is treated as a product requirement from the start, not a final compliance exercise. [6]

Lesson: Start with a measurable human outcome, make uncertainty visible, and build evaluation, review, and regulatory constraints into the workflow from day one.

Career Corner

Choose managers differently when choosing companies versus teams

Shreyas Doshi argues that “only follow great managers” is incomplete advice for experienced PMs joining fast-growing companies: those firms often lack mature management benches, even when the company itself offers substantial career upside. [8] In that context, avoid toxic or counterproductive managers, but be more tolerant of capable, overloaded, or less hands-on leaders. [8] After committing to stay at a company for several years, however, be far more selective when switching teams. [8]

For senior PMs, assess a prospective manager with four questions: Will they make it easier to get work done? Can you delegate meaningful work to them? Can they recognize top talent? Do they have organizational credibility to advocate for you? [8]

Tools & Resources

A shared, AI-readable planning repository

A practical workflow for PM planning is to export planning materials into a folder structure in VS Code, organize them into areas such as customers, market analysis, and planning, then use Claude Code or Codex against that repository. [4] Store interviews alongside strategy so AI can cross-reference recurring themes during discovery. [4]

To make the repository useful across the team: keep clear READMEs and a dedicated plans folder, convert large documents into Markdown, review every AI-generated plan before accepting it, and use a second agent or thread to critique important outputs. [4] The reported benefit is faster design/prototyping and preserved institutional knowledge when team members change. [4]

Sources

1. X post by @sachinrekhi
2. X post by @deanbloembergen
3. OpenAI’s models went rogue, and Google ships three Gemini models (but not the one you wanted)
4. Accelerating PM Workflows and Elevating Customer Value with AI | ProductTank London
5. How Photoroom Trained Themselves To Dream Bigger
6. Building AI for Women’s Health: How Fertility Combined Bayesian Diagnosis and Scan Automation
7. X post by @ttorres
8. The Great Manager Lie