

SAHIL SAPRA

Charlotte, NC 28273 | (216) 848-7236 | sahilsapra391@gmail.com | sahilsapra.com | [LinkedIn](#) | [GitHub](#)

August 20, 2026

[Plaid Hiring Team](#)

Re: Staff Product Manager, Internal AI

The hard part of this job is not choosing the tool. It is the week after the rollout, when the people you built it for quietly go back to the spreadsheet. I have been on the wrong side of that. At CAMP Systems I drove enterprise-wide adoption of LaunchDarkly to decouple deployment from release, and the first internal version stalled: I had built it around my own mental model of the workflow, and the feedback I got was that I had shipped my tool, not their tool. So I sat with people while they worked, rebuilt the rollout around the friction I could actually see, and adoption followed. Feature release frequency ended up 400% higher. The standing lesson I took from it is that no internal tool ships until I have watched a real user attempt the real task.

The AI half has been my own hands rather than a spec I handed off. At CAMP I built an internal test case generator grounded in retrieval-augmented generation over roughly 2,000 of our own ticket-to-test conversions, using dense embeddings against a vector index so the output followed our house conventions instead of generic ones and got better with every example added, no retraining. The generation was the easy part. The work was defining what good output meant against a house standard, building the human review loop that fed corrections back in, and iterating prompts and retrieval until people trusted it. Alongside it I directed the automated testing strategy that moved coverage from 62% to 81% without giving up release stability. I also built a LaunchDarkly API web app end to end in Cursor because the vendor dashboard was flag-centric while our workflow was release-centric; it consolidated seats and cut team OPEX 7%. Nobody asked for that one.

Building for people who do not write code is the part I would enjoy most, and I have done a version of it. The CEO of our second-largest account at World Kinect was getting 25 separate weekly revenue emails, one per airport location. Engineering had no bandwidth, so I built the replacement: a dynamic PowerBI digest with cumulative revenue and per-location drill-down, delivered every Sunday, with every figure reconciled against raw Snowflake data before it reached a CEO inbox. Twenty-five messy emails became one verified digest at zero engineering cost. On the founder side, I started SpecHawk, an agentic QA platform, secured four enterprise design partners, and raised pre-seed capital, and I measured those partners on real scheduled usage rather than demo enthusiasm, which is the same instinct this role needs when a shiny pilot has to survive contact with a quarter-end close.

Two things stated plainly rather than left for you to find. I am at roughly four years of product and analyst experience against your five-year floor, so I am under the bar you wrote, and you should weigh that against what the last two years actually contained. And I have never owned an internal tooling portfolio for Sales, Finance or HR specifically; my rollouts have been to engineering, QA and executive stakeholders, so the back-office workflows would be new territory I would have to learn from your teams. What I bring on day one is a real AI build record, an honest scar from a rollout that failed on adoption before it worked, and a Quantitative Economics degree behind the measurement conversations. I work in the United States without sponsorship, now and in the future, and I am willing to relocate to the Bay Area for this role. Thank you for your consideration.

Sincerely,

[Sahil Sapra](#)