

# SAHIL SAPRA

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

August 19, 2026

[Humana Hiring Team](#)

Re: Senior AI Technical Product Manager (R-425574)

I am applying for the Senior AI Technical Product Manager role. The line in your posting about building your own working software in Claude Code and Cursor is the reason I read the rest of it. That is how I already work. At CAMP Systems, LaunchDarkly's dashboard was flag-centric while our release workflow was not, and the seats were expensive, so I built the LaunchDarkly API web application myself in Cursor rather than filing a request. It consolidated seats and cut team OPEX 7%. Nobody asked for it.

The larger part of your posting is evaluation quality, and that is what I have spent the last two years on. At CAMP I owned the quality framework for a next-generation aviation platform and drove automated test coverage from 62% to 81% while holding enterprise release stability. The mechanism matters more than the number: I built an AI test case generator grounded in retrieval-augmented generation over roughly 2,000 of our own ticket-to-test conversions, dense embeddings on an Upstash vector index, so generated cases followed our house conventions and improved with every example added, with no retraining. I can explain RAG, embeddings and prompt iteration to a stakeholder and to an engineer without changing the story between the two. Then I left to build a company around exactly this problem. SpecHawk is a deploy-aware no-code agentic QA platform that maps a staging environment, generates and executes test specs against it, validates the Jira ticket and hands back screenshot evidence. Four enterprise design partners signed on and I raised \$100K pre-seed on a SAFE at a \$1M cap. Running weekly loops with those partners taught me to measure real scheduled usage rather than demo enthusiasm, which is the same instinct your posting is describing when it asks who decides what ships.

On business discovery and delivery at enterprise scale: onboarding one customer onto the next-generation platform took over two weeks of manual ETL, and I owned a 75-customer migration program, which is roughly three years done sequentially. I mapped the workflow end to end and automated the data injection and mapping instead of asking for headcount, taking capacity to three customers per day. All 75 migrated with zero operational downtime and ARR grew 104% year over year, two years running. When enterprise customers hit timeouts and the consensus called it a capacity problem, I traced request-level timing and found synchronous retrieval paths with no caching. Rebuilding it as domain-driven microservices cut data retrieval latency 68% and cost less than the infrastructure everyone wanted to buy.

Two gaps, stated plainly rather than left for you to find. My degree is a BS in Quantitative Economics with a co-major in Analytics and a minor in Computer Science, not the CS or Engineering degree you list, though I write Python and TypeScript and ship production code. And I have no healthcare or regulated-industry background, so your enterprise AI governance stage gates are something I would be learning rather than bringing. I am in Charlotte today and willing to relocate to a commutable distance of whichever office fits the team. I work in the United States without sponsorship now or in the future. Thank you for your consideration.

Sincerely,

[Sahil Sapra](#)