

SAHIL SAPRA

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

August 19, 2026

[Zendesk Hiring Team](#)

Re: Senior Product Manager, AI Agents — Berlin

I am applying for the Senior Product Manager, AI Agents role. Agentic AI is what I build rather than what I read about. At SpecHawk I founded and own a deploy-aware no-code agentic QA platform where AI agents map a staging environment, generate and execute test cases, validate Jira tickets with one click and capture screenshot evidence on a schedule. That was zero to one from an empty repository: I scoped the MVP to the smallest slice that produced a trustworthy test run on a design partner's real application, then ran a six-month roadmap through V1 into V2 and V3. Four enterprise design partners signed on, and I raised \$100K in pre-seed on a SAFE to fund it. Four years of B2B SaaS product management sit behind that.

The part of this role I find most interesting is the controls layer, because I have already learned the hard lesson about it. At CAMP I drove enterprise-wide adoption of LaunchDarkly and pushed feature release frequency up 400%, then built a governance web application on its API so the team had centralized control over releases in one place, which also cut team OPEX 7%. The first rollout stalled. I had built it around my own mental model, and the feedback was blunt: I had shipped my tool, not their tool. I sat with real users doing the real task, rebuilt around their actual friction, and adoption followed. Controls over how a system behaves are only worth what a customer can actually configure without help, and that is the standard I would hold the work to.

On grounding model output rather than hoping for it: I built an AI test case generator retrieval-augmented over roughly 2,000 of our own ticket-to-test conversions in an Upstash dense embedding index, so what it produced followed our house conventions and got better with every example added instead of needing a retrain. I can explain RAG, embeddings and prompt iteration to an engineer and then to a business stakeholder in the same meeting. It was part of taking test coverage from 62% to 81% while holding enterprise-grade release stability across a 75-customer migration that finished with zero operational downtime and 104% YoY ARR growth two years running.

Two things you should know before we spend time on this. I have not worked on a customer-service product, so the domain would be new to me even though the agent and governance work is not. I am also based in Charlotte, NC, so this would be a relocation to Berlin, which I am willing to do, and I would need German work authorization to be sponsored. If both of those are workable on your side, I would welcome the conversation. Thank you for your consideration.

Sincerely,

Sahil Sapra