

SAHIL SAPRA

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

TECHNICAL PRODUCT MANAGER | AI-ENABLED SAAS PRODUCT MANAGEMENT

August 19, 2026

[Momentive Software Hiring Team](#)

I am applying for the AI-Native Senior Product Manager role on the Agentic AI platform (R506). The line in the posting I want to answer directly is the one asking how AI has changed the way I work, what it has replaced, and where I have found its limits. I have a fairly specific view of all three, because I put these tools inside the product work rather than around it.

What it replaced first was the blank page. Requirements, acceptance criteria and comparative analyses now start from a generated first pass that I edit instead of write, and the edit is where the judgment sits. What it replaced second was arguing for a build. At CAMP Systems our release process was release-centric while the LaunchDarkly dashboard was flag-centric, and the seats were expensive, so I developed and deployed a LaunchDarkly API based web application end to end with Cursor. Nobody asked for it. It consolidated seats and cut team OPEX 7%, and building it was cheaper than making the case for it.

The limit I keep hitting is grounding. A general model writes plausible requirements and plausible test cases, which is not the same as correct ones. So I grounded our test generation model in RAG over roughly 2,000 of our own ticket-to-test conversions, using dense embeddings and a vector index, so output followed house conventions and got better with each added example without retraining. That was part of moving test coverage from 62% to 81% while holding enterprise release stability. The other limit is that AI is useless on the question of what to build, which stays a conversation with customers.

The product ownership half is what I spent two years on at CAMP. I owned the TFBO Next-Gen product area: comparative research against the legacy system, the roadmap that came out of it, a feature parity program translating workflow gaps into prioritized requirements and migration sequencing, and the move of 75 enterprise customers off competing platforms. Onboarding one customer took over two weeks of manual ETL, which made 75 of them mathematically impossible, so I drove automation of the data injection rather than ask for headcount and took capacity to three customers a day. All 75 migrated with zero operational downtime on a forced platform change, and ARR grew 104% year over year, two years running.

The closest thing I have to your job description is what I am doing now. At SpecHawk I own an agentic QA platform end to end: what jobs the agents perform, how a run decomposes into steps, and where a human has to sign off before anything touches a customer's environment. It has four enterprise design partners and \$100K pre-seed on a SAFE at a \$1M cap, and I run weekly design-partner loops that measure real scheduled usage rather than demo enthusiasm. I have led through influence rather than authority throughout, most visibly driving org-wide adoption of LaunchDarkly that raised feature release frequency 400%, after a first rollout that stalled because I had built it around my own mental model of the workflow instead of the users'. One thing you should hear from me rather than find out later: my domain is aviation B2B SaaS, not nonprofits and associations, so I would want to spend early time on customer conversations before I put a roadmap in front of anyone.

Sincerely,
[Sahil Sapra](#)