

SAHIL SAPRA

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

August 20, 2026

[Enterprise Health Hiring Team](#)

Re: Product Manager, Enterprise Health

Workflow-heavy enterprise product is the work I have actually been doing, and it is what I would want to be judged on here. At CAMP Systems, onboarding one enterprise customer onto our next-generation platform took over two weeks of manual ETL, and I owned a migration of 75 customers off competing platforms. Run sequentially that is roughly three years, which is not a schedule, it is a wall. Instead of asking for headcount I mapped the onboarding workflow end to end, found the data-injection step that was eating the two weeks, and drove the automation behind it until capacity moved to three customers a day. All 75 moved with zero operational downtime and ARR grew 104% year over year across two consecutive years. Most of that program was not engineering. It was telling leadership and customers, early, which feature gaps would break their live operations if I moved them too soon, and re-sequencing around it.

The AI part of this role has been my own job rather than something I scoped for someone else. At CAMP I looked at where AI could actually be applied to a workflow people were doing by hand, and built a 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 output followed our house conventions instead of generic ones and improved as the corpus grew, with no retraining. The generation was the easy half. The real 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 it held. Alongside it I directed the automated testing strategy that took coverage from 62% to 81% without giving up enterprise release stability.

Discovery through launch, then whether anyone actually uses it, is the part I have been burned by and learned from. I led enterprise-wide adoption of LaunchDarkly and lifted feature release frequency 400%, but the first rollout stalled, and the feedback was blunt: I had shipped my tool, not theirs. I rebuilt it around what I saw real users struggle with, and adoption followed. I run the same loop at SpecHawk, the agentic QA platform I founded, where four enterprise design partners are measured on real scheduled usage rather than demo enthusiasm.

One gap, stated plainly rather than left for you to find. I have not worked in health software, so medical surveillance and the rules that govern it are a domain I would be learning from your team rather than bringing on day one. What I do bring is four years of enterprise product where the constraint was an outside standard rather than my own taste, aviation maintenance and ERP rather than health, plus the LLM and retrieval work, SQL and Git in daily use, and prototyping in Claude Code and Cursor. I work in the United States without sponsorship, now and in the future. Thank you for your consideration.

Sincerely,

[Sahil Sapra](#)