

SAHIL SAPRA

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

August 19, 2026

[Klaviyo Hiring Team](#)

Re: Senior Product Manager, AI Analytics

Trust is the first of your five pillars, and it is the one I would want to be judged on. At World Kinect the CEO of a top-two account was receiving 25 separate weekly revenue emails, one from each airport location. Engineering had no bandwidth, so I built the replacement myself: a dynamic PowerBI digest with cumulative revenue and per-location drill-down, delivered every Sunday. The part that took the time was not the build. It was checking every figure against raw Snowflake data before any of it reached a CEO inbox, because a dashboard that disagrees with the system of record is worse than no dashboard. That is the same problem as an agent answering a marketer's performance question and having to agree with what the app shows.

The AI half of this job has literally been my job, not a spec I handed off. At CAMP Systems I 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 rather than generic ones and improved as the corpus grew, with no retraining. Most of that work was not generation. It was judging output quality against a house standard and iterating prompts and retrieval until it held. Alongside it I directed the automated testing strategy that took test coverage from 62% to 81% without giving up release stability. I then founded SpecHawk, an agentic QA platform where the product is an agent that discovers a staging environment, orchestrates its own test execution and validates the result, now with four design partners and a pre-seed raised.

The query engine migration in this posting is familiar ground too. I owned the product side of moving a legacy monolith to a domain-driven microservices architecture, and when consensus said the enterprise timeouts needed more infrastructure I traced request-level timing instead and found synchronous retrieval paths with no caching. Targeted caching and lazy loading on the hot paths cut data retrieval latency 68% and the timeouts stopped. Separately I mapped an onboarding workflow to the manual ETL step that cost two weeks per customer and automated the data injection behind it, moving capacity to three customers a day, which is what made a 75-customer migration arithmetically possible at all.

Two gaps, stated plainly rather than left for you to find. I have never owned an analytics or BI product as the product, so metric consistency and steady-state reporting health are things I have cared about from the reporting seat and the quality seat, not from the roadmap. And martech and e-commerce are new to me; I come from aviation B2B SaaS, so your customers' workflows are something I would be learning from your team rather than bringing on day one. What I bring on day one is the AI product quality work, the data grounding, and the habit of not letting a number ship until it reconciles. I work in the United States without sponsorship, now and in the future, and I am willing to relocate for this role. Thank you for your consideration.

Sincerely,

[Sahil Sapra](#)