

SAHIL SAPRA

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

August 19, 2026

Forcepoint Hiring Team

Re: Senior Product Manager, AI & Data Security

Your posting describes a lifecycle: discover, classify, assess, remediate. The end of it I have actually lived is assess, and the trust problem that sits under it. At World Kinect the CEO of our second-largest account was receiving 25 separate weekly revenue emails, one from each of his 25 airport locations. 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 build was the easy half. The real work was reconciling every figure against raw Snowflake data before any of it reached a CEO inbox, because a number that disagrees with the system of record is worse than no number at all. A finding a security team is about to act on has exactly that property, and a false positive costs them more than silence would have.

The AI half of this role has been my own job rather than a spec I handed to someone else. 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 instead of generic ones and improved as the corpus grew, with no retraining. Most of that work was not the generation. It was deciding what good output meant against a house standard, building the human review loop that fed corrections back in, and iterating retrieval and prompts until it held. Classification is older ground: at Keiretsu Forum I built Python classification models over roughly 250 early-stage companies to automate a due diligence pass. Alongside the generator I directed the automated testing strategy that moved coverage from 62% to 81% without giving up release stability, which is the same breadth-against-confidence call a classification engine makes every day.

You say the role is execution-biased, so here is execution. Enterprise customers were hitting timeouts on our legacy monolith and the consensus fix was more infrastructure. I traced request-level timing, found synchronous retrieval paths and no caching, and owned the product side of a move to domain-driven microservices with reusable REST APIs, caching and lazy loading aimed at the hot paths. Data retrieval latency fell 68%, the timeouts stopped, and it cost less than the hardware would have. Separately, onboarding one enterprise customer took over two weeks of manual ETL while I owned a 75-customer migration, which run sequentially is about three years. I drove automation of the data injection instead of asking for headcount, took capacity to three customers a day, and all 75 moved with zero downtime while ARR grew 104% year over year across two consecutive years.

The gap, stated plainly rather than left for you to find. I have not built DSPM, DLP or CASB products and I have not worked inside a cybersecurity vendor, so the domain vocabulary and the competitive field are things I would be learning from your team in the first quarter rather than bringing on day one. What I do bring is enterprise product management with real engineering depth, the retrieval and LLM work above done with my own hands, a security-adjacent build in SpecHawk's Security Hub for dependency risk analysis, AWS Certified Cloud Practitioner behind the cloud conversations, and the habit of not letting a number ship until it reconciles against the source. I work in the United States without sponsorship, now and in the future, and I am willing to relocate to Virginia for this role. Thank you for your consideration.

Sincerely,

Sahil Sapra