

SAHIL SAPRA

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

PRODUCT MANAGER | RISK DECISIONING, PRODUCT STRATEGY & PLATFORM GROWTH

August 20, 2026

Chime Hiring Team

I am applying for the Product Manager, Transaction Risk Decisioning role. The part of the posting I keep coming back to is the framing of risk as a business tradeoff rather than a loss number. That is the argument I have had to make repeatedly, and usually against the room.

At CAMP Systems I directed the automated testing strategy that took coverage from 62% to 81%. The interesting decision there was not the number, it was trading breadth of coverage against confidence in each individual result, and being able to say out loud which one we were buying and what it cost. Separately, enterprise customers were timing out on our legacy platform and the consensus fix was more infrastructure. I traced request-level timing instead, found synchronous retrieval paths and no caching, and led the move to domain-driven microservices with reusable REST APIs, caching and lazy loading aimed at the hot paths. Data retrieval latency dropped 68%, the timeouts stopped, and it cost less than the infrastructure everyone wanted to buy.

Trust under a bad outcome is the other half of this job, and the closest thing I have to it is the 75-customer migration I owned. Some of those customers had feature gaps that would have broken live aviation operations if we moved them on schedule. I told leadership and the customers the specific gaps, ran a feature parity program, re-sequenced the migrations and flagged slips before anyone asked. All 75 moved with zero downtime and zero churn on a platform change they did not choose, and ARR grew 104% year over year two years running. People forgive a block far more readily when you tell them exactly what it is and what gets them through it.

On the systems side, I build rather than review. I founded SpecHawk, an agentic QA platform where the agent maps a staging environment, generates and executes test cases and returns screenshot evidence, and I grounded its test generation in RAG over roughly 2,000 of our own ticket-to-test conversions so the output followed house conventions. Earlier, I built Python classification models that scored around 250 early-stage startups for an in-house venture fund.

Where I am thinner is card network mechanics. I have not worked authorization reason codes or chargeback frameworks, and I would be learning that vocabulary from your risk strategists rather than bringing it. What I would bring on day one is the decisioning-under-dependencies half, and the habit of stating the tradeoff in revenue terms before anyone asks for it.

Sincerely,
Sahil Sapra