

SAHIL SAPRA

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

August 19, 2026

[Cohere Hiring Team](#)

Re: Product Manager, Platform Experience & Developer Product

Dear Cohere Hiring Team,

I am applying for the Product Manager, Platform Experience and Developer Product role. The part of your scope I have actually lived is the platform layer underneath the product. At CAMP Systems I led an enterprise aviation platform off a legacy monolith onto a domain-driven microservices architecture, with reusable RESTful APIs, caching and lazy loading aimed at the hot paths, which cut data retrieval latency by 68% and eliminated the timeouts our largest customers kept hitting. The consensus at the time was to buy more infrastructure. I traced request-level timing instead, showed it was an architecture problem rather than a capacity one, and the fix was both faster and cheaper. In the same stretch I moved 75 enterprise customers off competing platforms with zero operational downtime while ARR grew 104% year over year, two years running.

Honestly, what got my attention is the range packed into one job. Shared access to models on one end and dedicated single-tenant deployments on the other, synchronous real-time calls and high-volume async workloads sitting behind the same APIs, the same dashboards, the same developer experience. Those surfaces pull against each other and the developer is the one who sees the seams. I have built the small version of that problem. Our release control vendor's dashboard was flag-centric while our workflow was release-centric, so I wrote a LaunchDarkly API web application in Cursor that gave the teams centralized feature flag monitoring and control, consolidated seats and took 7% off team OPEX. The rollout of LaunchDarkly itself stalled at first because I had built it around my own mental model instead of the engineers', so I watched real users work, rebuilt it around their friction, and release frequency went up 400%.

I should name the gap rather than let you find it. You ask for five or more years and I am at four, and I have never owned model serving or an inference platform, so that layer I would be learning rather than arriving with. What I do bring is four dense years of enterprise B2B SaaS platform ownership, API and cloud infrastructure depth I can hold a real conversation with your engineers about, and a year running my own AI product: SpecHawk, a deploy-aware agentic QA platform I took from nothing to a V1 MVP, four design partners, a developer CLI and GitHub integration. I am in Charlotte, authorized to work in the US without sponsorship, and happy either remote or relocating for the San Francisco or New York option. Thank you for your consideration.

Sincerely,

[Sahil Sapra](#)