

Alex Chen

Austin, TX (US remote) | alex.chen@demo.example.com | [LinkedIn](#)

Cover Letter: Senior AI Platform Engineer — Aurora Labs (demo posting)

Aurora Labs Remote (US) 2026-07-11

Hi Aurora Labs team,

Your job description asked for someone who has actually taken LLMs from prototype to production traffic — not just fine-tuned notebooks — and who cares about eval and observability as much as raw model quality. That is exactly the arc I've been on for the past four years, and it is why I'm writing.

At TechFin Corp I led the ML platform team that grew from 2 models in production to 15+. I built the internal LLM serving stack (registry, A/B routing, feature store, per-tenant limits), designed the real-time fraud detection pipeline that runs at 99.7% precision and 50 ms p99, and stood up the drift + performance monitoring that lets us actually sleep at night. On the side I maintain LLM Eval Toolkit (open source, ~600 GitHub stars) — the same eval philosophy your JD calls out.

- **Cut model deploy time from 2 weeks to 4 hours**, by wiring CI/CD with automated eval gates against a held-out set — a direct mirror of your "model rollout with confidence" bullet.
- **Shipped LLM reranking on top of a collaborative-filtering system**, with 18% conversion uplift over 12 A/B weeks, so I've felt what "production LLM traffic" actually behaves like.
- **Ran production ML on-call for 3 years**, and turned that pain into the drift-monitoring + retraining triggers that dropped page volume ~35%.

The three problems I'd want to help Aurora Labs pick up first: (1) turning your eval suite into a rollout gate rather than an offline report; (2) building a clean seam between the model layer and the product surface so experimentation stops leaking into product code; (3) making cost per successful request a first-class SLI alongside latency and quality.

I'd love a 30-minute conversation to make sure this maps to real problems on your side, not just my reading of the JD. Happy to walk through a concrete rollout-gate design I built at TechFin, or dig into whatever specific piece of your platform is most on fire this quarter.

— Alex

* This is a demo cover letter generated for the Clawvard career-ops-copilot course showcase. Alex Chen and Aurora Labs are synthetic demo entities; no real candidates or employers are referenced.