Kartik Aslia

Last updated in September 2026

RL Environments Engineer | Agentic Task & Environment Pipelines | Backend & Infra

Gurgaon, Haryana

asliakartik@gmail.com +91 77019 23367 github.com/Skizzy-create

Summary

RL environments engineer who builds the pipelines and tooling that turn environment ideas into validated agentic coding tasks at scale. Recently built — solo, in three days — a parallel Docker-in-Docker capture farm now used across the team, cutting a 50-task batch from ∼150h to ∼4h, and personally shipped **36 of 50 tasks (72%)** in a client delivery — the highest volume on the team. Authored the delivery-QC gates that catch reward hacking and grader loopholes to hold quality as volume grows. Currently an expert reviewer hardening complex agentic environments and validating human/LLM trajectory annotations for research-and-planning agents. Backed by production backend experience — root-cause analysis, sandboxes, CI, and platforms scaled to 1000+ concurrent users — and a track record of owning delivery under deadline.

Skills

- **RL Environments & Eval:** RL environment generation, agentic coding-task authoring, automated grading & verification, reward-hacking / grader-loophole defense, trajectory-annotation review (human & LLM), pass@1 calibration, failure analysis
- **Platform & Infra:** Docker / Docker-in-Docker, cgroups & network namespaces, sandboxes / code-execution infra, Grafana LGTM stack (Loki, Tempo, Mimir), Parquet / pyarrow, parallel capture farms, AWS (EC2, ECS, RDS, S3, CloudFormation)
- **Languages:** TypeScript, JavaScript, Python, SQL, C++
- **Backend & Data:** Node.js, Express, PostgreSQL, MongoDB, Prisma ORM, Redis, REST APIs, WebSockets, Server-Sent Events, worker threads
- **Applied AI:** LLM applications, RAG / Self-RAG, LangChain, LangGraph, prompt engineering, semantic search, OCR / document AI
- **Delivery & Ownership:** Delivery ownership, QC gate design, contributor review, throughput optimization, cross-functional coordination, prioritization under deadline

Experience

Bespoke Labs — *RL Environments Engineer – Expert (Contractor, Remote)* *April 2026 – Present*

Pod Lead → Senior Reviewer (Contractor) — Nebula Aurora *Apr 2026 – Jul 2026*

- Joined as **Pod Lead**: authored and quality-checked agentic RL-environment coding tasks (Pod 6) against a strict client data contract; coordinated pod contributors, onboarded new members, and cleared blockers to keep delivery moving.
- Promoted to **Senior / Task Reviewer** while owning delivery — delivered high-accuracy reviews and second opinions across the pod, enforcing evidence, correctness, and reachability standards before client hand-off.
- Recorded the **highest client-accepted task-delivery count** on the pod worksheet despite a **∼15–20 day** active window versus peers running 6+ months, while consistently taking on the most rigorous task classes.

RL Environments Engineer (Contractor) — Byzanite *Jul 2026 – Aug 2026*

- Built — solo, in **3 days** — a parallel **Docker-in-Docker** capture farm now used by the whole team, cutting a 50-task batch from **∼150h to ∼4h (≈37×)** and a 10-task batch **∼21h → ∼4h**; added a time-shift scheduling algorithm that keeps parallel captures delivery-compliant and cleared the scaling blockers (Docker address-pool exhaustion, a cgroup-delegation race, slot collisions).
- Built RL environments for **root-cause analysis** on real production codebases — incident worlds (Ethereum nonce/broadcast, Solana deposit-watermark, Redis/Postgres cascades, Kafka/JWT, KYC/AML) with retrievable evidence and graders calibrated to a target pass@1.
- Authored the **delivery-QC gates** that defend quality at scale — a gate that *executes* each cited query instead of trusting the text, a bare-repo git-mirror check (I25), a relationship-class enum check (I19), and a build-time internal-material leak guard — catching grader loopholes and taking a 10-task package to **0 errors / 0 warnings** on the client validator.
- Built a **Parquet-based telemetry slot-shift and merge** (pyarrow over Loki / Tempo / Mimir) so multiple incidents serve from a single stack over disjoint time windows, with fail-closed collision checks.
- Parachuted in **∼2 working days before deadline** (onboarded Fri/Sat for a Tuesday delivery) to rescue the pilot: briefed and shipping environment fixes within hours, landed the 10-task client pilot, and after spec iterations delivered the 50-task package.
- Personally shipped **36 of 50 tasks (72%)** in the final delivery — **top of five contributors** (next-highest 10; median ∼2), the highest personal volume on the team — top contributor across **214 commits**, including a salvage that recovered **19/19** captured environments with zero recapture.

Expert Reviewer (Contractor) — AlphaRecon *Aug 2026 – Present*

- Serve as **expert reviewer** on complex agentic coding environments — reviewing each task for prompt quality, environment complexity, and up-to-date research/planning context, so the environments train agents to **research, plan, and review** more effectively.
- Review **trajectory annotations** — validating human- and LLM-produced trajectory labels for correctness before they enter the training and evaluation set.

Director of Technology & Product Lead *Feb 2025 – Present*

KarmicDD, CLAR AI Workspace, CLAR for Microsoft Word

- Owned backend system delivery for a compliance-heavy due-diligence platform end-to-end — architecture, development, and deployment — coordinating with product, legal compliance, and customer success to translate regulatory requirements into engineering specs under audit timelines.
- Scaled the platform to **1000+ concurrent users** with performance gains (**200ms → 50ms**) via caching, async task queues, and worker-based processing; built backend systems in **Node.js / TypeScript** with **PostgreSQL**, **MongoDB**, **Prisma**, and **JWT / CSRF** on client-isolated storage.
- Led applied-AI workflows for document generation, legal analysis, and research using **Gemini**, approval gates, and audit logging; shipped document pipelines for **PDF, DOCX, XLSX, PPTX, and OCR** with real-time **SSE** and **WebSocket** flows for streaming chat and long-running AI tasks.

Software Engineer (Contract) *Dec 2024 – Jan 2025*

Autonomous Web3 / Quant Intelligence Engine

- Built an agentic market-intelligence engine combining on-chain signals, news feeds, and **147+ technical indicators** to support automated portfolio decisions and risk-aware execution.
- Led architecture and system-design decisions independently, interfacing directly with the founding team on product requirements and technical trade-offs.

Projects

Encryption App – *React, Node.js, AES-256, JWT* *GitHub Live*

- Built a secure encryption and decryption web app with JWT authentication, AES-256 encryption, QR-code generation, and multi-algorithm secure sharing flows.

Paytm Clone – Full-Stack Payments App – *React, Express, MongoDB, JWT* *GitHub Live*

- Developed a full-stack payments application with secure authentication, transaction management, balance monitoring, and fund transfers using React, Express, and MongoDB.

Education

B.E., Computer Science – Chandigarh University, Punjab *CGPA: 7.98/10* *2022 – 2026 (Expected)*

Achievements

- Led team to **2nd place** at TON Gurugram BootCamp, winning **1500 USD in USDT**.
- Achieved **All India Rank 5** at **IDEATHON-2023**, competing against **500+ teams**.
- Led team to **Top 6** at **Smart India Hackathon (SIH)**.