DevOps / SRE Lead

Team: Engineering **Reports to:** [CTO / Head of Engineering] **Location:** Gurgaon, on-site **Hours:** Overlap with US dealership hours until [X] am IST, remote after [Y] pm. On-call is shared across the team on a published rotation. **Experience:** [X]+ years in infrastructure or SRE, [Y]+ running production on-call **Team today:** [You start as the first dedicated DevOps hire / you inherit N engineers]

What Spyne does

Spyne builds AI for US car dealerships: **Vini AI** agents that answer their phone calls, texts and leads, and **Studio AI**, which turns photos of cars into listing-ready images and video. When our platform is slow or down, a rooftop's listings and phones are affected during its selling day.

The role

You lead DevOps and SRE at Spyne. You own reliability, deployment, observability and infrastructure cost across all products. Product teams own their services; you own the platform, the standards, and the incident process they run on. Our production is not a plain web stack: Studio runs GPU and media-processing pipelines on thousands of photos and videos, and Vini runs real-time voice, where latency is the product. You are a builder first: you use agentic tools to automate the work others do by hand, and you set the standard for how the team runs production.

What you will do

- Own uptime, incident response and root cause analysis for production, including vendor failover and degradation paths for third-party telephony and model APIs, and a fair, sustainable on-call model.
- Build and run the deploy pipeline, environments and rollback paths.
- Own platform security: secrets, access control, and [compliance scope].
- Own observability: logs, metrics, alerts and dashboards that catch problems before a customer or CSM does.
- Own cloud cost: a published plan, a dashboard that shows progress, and savings that stick as the rooftop count grows.
- Automate toil with agents and scripts; build runbooks and let the rest of engineering self-serve.
- Build and lead a small DevOps/SRE team over time; hire and raise the bar.

What we are looking for

- You have run production for a product customers pay for, and led or been on the hook for its on-call.
- Deep infrastructure knowledge: AWS, containers and orchestration, CI/CD, infrastructure as code, networking, databases in production.
- Strong agentic builder: you automate with Claude Code or similar tools and can show examples.
- You stay on an outage until the cause is proven and prevented, and you tell people early when something is going wrong.
- Proven incident leadership: calm under pressure, clear written post-mortems, fixes that remove the class of problem.
- Cost discipline: you have cut cloud spend without hurting reliability.
- Security and secrets management in production.
- Nice to have: GPU or media-processing workloads, voice or real-time systems.

How to apply

Email **[apply address]** with:

1. An incident you led: what broke, how you found the cause, and what changed so it could not happen again.
2. Two or three sentences on something you automated that others used to do by hand.

No cover letter, no take-home.

How we will hire

1. Intro call: scope, ownership and hours.
2. Technical deep dive on infrastructure you have run and the incident you submitted.
3. Practical exercise (about 2 hours, AI allowed): diagnose a failing service from logs and metrics, then propose the fix and the prevention.
4. Conversations with engineering leadership.

What success looks like

- Published SLOs for Vini and Studio, met each month: [uptime], [voice latency p95], [Studio turnaround time].
- Time to root cause and time to recovery tracked and falling.
- Cloud cost per rooftop falling against a published plan, then held as we grow.
- A deploy and on-call process the team trusts.

Benefits

- **Pay:** market-leading salary and stock options.
- **Health and life cover:** group medical (GMC), personal accident (GPA) and term life (GTLI) insurance for you and your family.
- **Learning:** LinkedIn Learning access and mentorship programs.