

Akshit Pareta

Bundi, Rajasthan, India • akshitpareta01@gmail.com • +91-9680810188

akshitpareta.com • github.com/akshitpareta • linkedin.com/in/akshitpareta

Re: Application for Forward Deployed Engineer

Dear Hiring Team,

For the past four years I have done exactly what a Forward Deployed Engineer does — walk into a business that cannot fully articulate its problem, figure out what is actually broken, and ship a system that stays in production. I have done this 50+ times across eight industries. I am applying to [COMPANY] because [SPECIFIC_THING] — for example, your work deploying AI into regulated enterprises mirrors a multi-tenant legal RAG platform I shipped for a six-founder law firm, used daily via WhatsApp, saving each attorney 15–30 hours a week and lifting case capacity 50%.

What I bring that most FDE candidates do not:

- **Production AI infrastructure ownership** — I built and operate a self-hosted 2× RTX 5090 cluster (64GB VRAM) running quantized Qwen-72B (INT4) via vLLM with QLoRA fine-tuning pipelines. Zero cloud GPU cost; ~Rs.40K/month avoided.
- **Multi-agent orchestration at scale** — Swarm OS coordinates 50+ concurrent LLM agents across five frontier models with per-agent context budgeting, max-iteration caps, and dependency-cycle detection so it runs unattended, not just as a demo.
- **Customer-facing shipping velocity** — 28 production deployments and 2 mobile apps (5,000+ users, Rs.50L revenue), each shipped end-to-end: discovery, architecture, build, deploy, iterate.

I am commerce-educated and self-taught as an engineer. For an FDE that is not a weakness — it means I read the business and the P&L before I write the code, so my deployments get adopted because they solve the real problem, not the stated one.

I would love to walk you through the legal RAG architecture — retrieval design, the failures I debugged, and the 340-case eval harness that gates every deploy. Worth a 20-minute conversation?

Best regards,

Akshit Pareta

akshitpareta.com • +91-9680810188