



Senior / Staff Software Engineer (Track to Lead Engineer)

Austin · Onsite preferred (hybrid in rare cases) · Competitive base plus early-stage equity

Reflekt Legal is the AI workforce for law firms. Firms hire our AI employees, who work the same channels their teams already use with humans (phone, email, SMS) and plug into the case management systems firms already run on. Today they handle client intake end to end. We're growing fast and approaching an inflection point, and we need an exceptional engineer to help us get there. Our goal is simple: let lawyers be lawyers, and help firms deliver client service that feels above and beyond.

Why this role

You own the stack and set technical direction alongside the founders. You have a real say in the big architectural decisions, and what you ship reaches real users quickly. Perform well and you're on track to Lead Engineer.

What you'll do

- Own features end to end across frontend, backend, and the infrastructure they run on.
- Build and scale our AI employees platform: the application, the FastAPI services, and the data flows that connect them.
- Run our cloud and telephony backbone on GCP and Kubernetes, keeping production reliable and low-latency as we grow.
- Lead our real-time voice work and own its hardest challenges: pushing the voice to be the most realistic, driving down latency, and expanding what our AI employee platform can do.
- Set technical direction with the founders and raise the engineering bar as the team grows.

Who you are

You own systems end to end, you have deep cloud infrastructure experience at the staff level, and you use AI in your own work every day. You don't need deep voice, telecom, and applied AI all on day one.

Our stack is Python and FastAPI on the backend, GCP and Kubernetes for infrastructure, a modern TypeScript frontend, and real-time media on top. You should be comfortable across most of it and strong in at least one layer.

Ideal profiles:

- Staff or senior full-stack engineer with 4+ years shipping production systems, who has owned cloud infrastructure at scale, ideally with voice, SMS, or real-time media experience.
- Engineer building with applied AI beyond prompt engineering (speech models, ML pipelines, real-time inference) who wants to do it full time.
- Engineer from a voice, telephony, or legal-tech company who already knows the real-time stack and the integration landscape.

Nice to have

Experience at a voice AI, speech, or telephony company (TTS, STT, SIP, VoIP, latency optimization), or running Kubernetes in production at an early-stage company.

How to apply

Email careers@reflektlegal.ai with a short note on why this role interests you and the hardest technical problem you've shipped, ideally in voice, speech, or real-time systems. Our CTO reviews every application personally.