

Founding AI Engineer (Voice AI)

Bangalore · Full-time · In-office · 2–8 years experience

Build prompt-led voice agents and the real-time systems that keep them reliable in production.

ABOUT THE ROLE

AuroraX is building Google for the offline world. Most commerce still happens in local stores and services with no live pricing, availability, or inventory online. Our consumer voice AI calls local sellers on a buyer's behalf, gets real answers, and returns structured, verified results. An invite-only beta is live, with the US as our first market.

We are a stealth-stage, pre-seed funded team, founded by ex-YC founders, with the founding team split across the US and India. You would own the real-time voice pipeline that is the core of the product — roughly 80% prompt and agent engineering, with the remaining 20% backend engineering.

WHAT YOU WILL OWN

- Aurora AI's primary source of data — voice agents.
- The entire voice agent pipeline — STT, LLM, TTS, turn detector, and everything between them.
- Evals as infrastructure: continuous evals on live calls, regression suites for every prompt and model change, and a failure taxonomy that tells us exactly where calls go wrong.
- Model choices and tradeoffs: picking and swapping providers across the stack based on quality, latency, and cost, with fine-tuning where it actually pays off.
- The quality loop: listen to real calls, find the failure modes, ship the fix, and prove it worked with evals rather than vibes.
- Production reliability for voice agents — monitoring, incident response, and continuous improvement after ship.

YOU ARE A FIT IF

- You have owned a real-time voice AI pipeline that served live users in production. Overall, you have 2 to 8 years of engineering experience.
- You know where the milliseconds go: you can profile and cut latency across ASR, inference, and TTS, and you have opinions about streaming versus batch at every stage.
- You are comfortable below the model layer as well: websockets, streaming audio, and telephony infrastructure.
- You debug from call recordings and logs, and you ship fixes in days.
- You have experience with Python and LiveKit.

NICE TO HAVE

- You have fine-tuned or custom-trained your own models and can accurately predict when that beats prompting a frontier model.
- You have deployed models in Azure Foundry, AWS Bedrock, or similar.

-
- Prior work on agent reliability, continuous model evaluations, or conversation quality scoring.
 - IIT/BITS/NIT graduate preferred.
-

PROBABLY NOT A FIT IF

- Your AI experience is text-only: chatbots or RAG apps with no real-time audio in production.
 - You come from ML research with no record of shipping to live users.
 - You have no experience with Python or backend engineering — voice agents involve lots of moving parts to achieve good latency and quality, so you need to understand the pipeline end to end to contribute meaningfully.
-

COMPENSATION

Competitive cash plus a meaningful founding-team ESOP grant. Final band depends on experience and evaluation.

HOW TO APPLY

Send a short note on a voice pipeline you have owned: what you built, the latency you hit, and how you evaluated it. Applications go to join@aurorax.co.