

Most infrastructure problems worth solving are either scientifically hard or engineered at scale, rarely both at once. We are the exception. Over the past 18 months, we have built a physics based model of the global freight network from scratch: 1,400 ports mapped at the berth level, 28 chokepoints, 140 carriers, 7,500 vessels, and 54 million customs manifests that give us per account, per-lane intelligence no competitor can replicate. The science is done. We are currently building the product layer: the APIs, the agent pipelines, the model serving infrastructure, and the enterprise interfaces that make this intelligence usable for the supply chain teams who need it. That is the engineering problem we are hiring for, and it is a genuinely hard one: building a production-grade agentic system that ingests real-time signals across a global freight network, runs inference at the right latency for operational decisions, and surfaces the right action to the right person at the right moment. If you want to define the technical architecture of a company from day one. Not inherit someone else's decisions but make the founding ones and you want to do it on top of a scientific foundation that took 18 months to build and can't be bought, we should talk.

## Head of Engineering / VP Engineering

Pref. location: Hyderabad Salary: Negotiable

### WHAT THE RIGHT PERSON LOOKS LIKE

- 8 - 14 years of experience, with meaningful time at a startup or high-growth tech company - not a long enterprise stack career
- Has shipped backend systems to production and led or been a core member of a small, high-output engineering team (3 - 10 people)
- Has seen how agentic AI infrastructure works - LLM orchestration, model serving, agent pipelines - even if they haven't architected one from scratch
- Comfortable across the full backend stack: APIs, cloud infrastructure (AWS/GCP/Azure), databases, and system design
- Strong opinions on code quality, shipping cadence, and building for reliability from day one
- Ideally a founding engineer from a prior startup, looking for their next mission - not someone seeking a stable senior role

### WHAT YOU'LL OWN

**Full product infrastructure:** the layer that makes Tau's ML models accessible, usable, and reliable for enterprise clients

**API architecture:** how the agent pipeline is exposed to customers' systems and integrated with their existing tools

Backend systems connecting our proprietary models to the front-end product

**Near-term:** shipping a working product that a customer can interact with - this is the immediate, concrete deliverable

**Over time:** engineering team build-out, hiring, technical roadmap, and culture

### WHAT YOU'LL TAKE AWAY FROM THIS EXPERIENCE

Tau has spent 18 months building physics-based models, deep learning based intelligence layers that no competitor can replicate. The person who builds the product on top of that foundation gets to say they helped turn raw scientific IP into the intelligence system global logistics runs on.

