



CTO / VP of Engineering (AI startup / San Francisco)

HITCONTRACT

hitcontract.lt

Vilnius, Lithuania

About the assignment:

Location

San Francisco,
Lithuania

Rate (after tax)

from €7000/Month

Duration

Full time position

Extension (project)

No

Remotely (optionally)

No

Expire On

2025-12-31 (15 hours
from now)

Description

About the Company

This AI startup is building an AI Sales Superintelligence — an agentic operating system that automates the entire go-to-market (GTM) motion for modern B2B companies.

Think ZoomInfo meets OpenAI: a unified AI platform where intelligent agents source, engage, and close deals autonomously across channels.

The mission is bold: enable companies to scale to \$1B ARR with fewer than 10 reps through vertical AI automation.

The startup is backed by Y Combinator (W25) and top-tier investors, led by an experienced team that has previously scaled startups from \$50M → \$600M+ ARR (Deel, Ramp, Vanta) and built enterprise-grade AI systems at scale.

They move fast, ship daily, and operate like a special-ops team — lean, high-impact, and obsessed with engineering excellence.

The Role

The company is hiring a Head of Engineering to guide its next phase of growth — evolving from a high-velocity early-stage team into a world-class engineering organization.

You'll work directly with the founders and own everything from the technical roadmap and architecture to team culture and execution.

This is a founding leadership role with significant impact, autonomy, and equity — the opportunity to shape both the product and the engineering culture that defines the startup's future.

What You'll Do

- Own the engineering roadmap and translate company vision into a clear, scalable technical strategy.
- Lead and mentor a high-performing team across product, infra, and AI; scale the team from 3 → 20+ engineers.
- Architect and scale a multi-agent system (React, Node.js, MongoDB, AWS, microservices, Chrome extensions).
- Guide the evolution of the AI stack (LLM pipelines, LangChain-style frameworks, fine-tuning, agent orchestration).
- Drive excellence in performance, reliability, and security — supporting enterprise-grade customers.
- Build lightweight but rigorous engineering processes (sprint planning, on-call rotations, code health).
- Partner with Product and Design to ship daily while preserving long-term architectural integrity.

Who You Are

- 8-12+ years of engineering experience, including 3+ years leading teams in high-growth startups.
- Proven record of taking products from 0 → 1 → scale, ideally in AI, SaaS, or data-heavy environments.
- Deep technical expertise in distributed systems, backend architecture, and modern web stacks.
- Experience with microservices, async data pipelines, and scalable cloud infrastructure.
- Strong leadership DNA — you build, mentor, and inspire high-agency teams.
- Passion for AI automation, LLMs, and emerging agent architectures.
- Thrive in a high-autonomy, high-velocity environment with clear ownership.

Bonus Points If You:

- Have shipped or led AI/LLM-powered products.
- Have experience with data, CRM, or GTM intelligence platforms (e.g., ZoomInfo-like systems).
- Understand enterprise-grade security and compliance (SOC2, ISO27001, etc.).
- Have worked with LangChain, vector databases, or agent orchestration frameworks.

Why Join

- Build the GTM Superintelligence and work at the frontier of applied AI for revenue.
- Massive ownership: define architecture, culture, and scaling strategy from the ground up.
- Work with elite founders with strong track records in AI and hypergrowth startups.
- Y Combinator-backed: strong funding, traction, and support.
- Hybrid flexibility: San Francisco HQ (5x/week) or remote-first for exceptional candidates.
- Competitive package: meaningful equity + salary + benefits (health, wellness, parental leave, unlimited PTO).

This is not an average leadership role — this is a starship seat.

If you're ready to push the boundaries of AI engineering and help build an iconic company, this is your launchpad.

Interview Process

1. **Intro Call (30 min)** — Meet a founder to discuss experience, leadership philosophy, and motivations.
2. **Technical Deep Dive (60 min)** — Review architectural decisions, scaling experience, and leadership approach.
3. **Founder Round (30-45 min)** — Final conversation focusing on product mindset, velocity, and culture fit.
4. **Offer** — If it's a match, the process moves quickly.

About the AI Startup

This company is redefining how enterprises sell — replacing fragmented sales stacks with vertical AI agents that automate the entire revenue cycle.

The platform unifies data enrichment, engagement, and pipeline intelligence into a single AI system, enabling

teams to focus on what humans do best: relationships and strategy.

They are building the infrastructure for the AI Sales Era — where human and machine collaboration becomes the new GTM superpower.

Required Skills

DATA BASE

Mongo DB +6 years

TECHNOLOGY

Amazon Web Services +6 years

JAVASCRIPT

React.js +6 years TypeScript +6 years Node.js +6 years