



## Senior Software Engineer – AI Infrastructure

### Location

United Kingdom; Brazil; Canada; Cyprus; Czech Republic; Germany; Ireland; Lithuania; Mexico; Poland; Portugal; Romania; Spain; Switzerland; United Arab Emirates

### Employment Type

Full time

### Location Type

Remote

### Department

Engineering ›

AI & Machine Learning

## OverviewApplication

### Building the Future of Crypto

Our Krakenites are a world-class team with crypto conviction, united by our desire to discover and unlock the potential of crypto and blockchain technology.

### What makes us different?

Kraken is a mission-focused company rooted in crypto values. As a Krakenite, you'll join us on our mission to accelerate the global adoption of crypto, so that everyone can achieve financial freedom and inclusion. For over a decade, Kraken's focus on our mission and crypto ethos has attracted many of the most talented crypto experts in the world.

Before you apply, please read the [Kraken Culture](#) page to learn more about our

internal culture, values, and mission. We also expect candidates to familiarize themselves with the Kraken app. Learn how to create a Kraken account [here](#).

As a fully remote company, we have Krakenites in 70+ countries who speak over 50 languages. Krakenites are industry pioneers who develop premium crypto products for experienced traders, institutions, and newcomers to the space. Kraken is committed to [industry-leading security](#), [crypto education](#), and [world-class client support](#) through our products like [Kraken Pro](#), [Desktop](#), [Wallet](#), and [Kraken Futures](#).

## **Become a Krakenite and build the future of crypto!**

### **Proof of work**

#### **The team**

The AI Infrastructure team builds and operates the production systems that power intelligent agents at scale. This team sits at the foundation of the agent platform, ensuring that model inference, orchestration, and execution layers are reliable, observable, and performant under real-world load.

Working closely with the Agent Systems team and broader infrastructure partners, this group owns the core primitives that enable agents to safely operate across internal systems. The environment is high-scale and high-stakes — systems serve millions of users and must meet strict reliability, latency, and correctness standards.

This is a deeply production-oriented team. Engineers here combine strong systems thinking with applied ML infrastructure experience, building in Rust and operating services where performance and failure modes matter.

#### **The opportunity**

- Design and build the infrastructure layer powering AI agent systems in production
- Develop high-performance Rust services that handle model inference, orchestration, and execution
- Architect scalable systems capable of supporting millions of users and high request throughput
- Build reliable ML infrastructure and MLOps patterns for model deployment, evaluation, and monitoring
-

- Define guardrails, observability, and failure handling for agent-driven workflows
- Optimize latency, throughput, and cost across inference and orchestration layers
- Partner closely with the Agent Systems team to translate experimental prototypes into hardened production systems
- Contribute to foundational infrastructure decisions in a high-scale, high-impact environment

### **Skills you should HODL**

- 5+ years of experience building and operating high-scale production systems
- Strong proficiency in Rust and systems-level programming
- Deep understanding of distributed systems, reliability engineering, and performance optimization
- Experience operating services serving millions of users or high-throughput workloads
- Familiarity with ML infrastructure, model serving, or MLOps in production environments
- Experience designing observability, monitoring, and failure recovery systems
- Strong collaboration skills working across infrastructure and applied engineering teams
- High ownership mindset in high-stakes production environment

### **Nice to have**

- Experience building infrastructure for agent-based or LLM-powered systems
- Background in high-performance networking, async systems, or low-latency architectures
- Experience with container orchestration and cloud-native infrastructure
- Familiarity with evaluation frameworks and model performance monitoring at scale
- Experience working in fast-moving 0→1 or platform-building teams

*Unless a specific application deadline is stated in the job posting, applications are accepted on an ongoing basis.*

*Please note, applicants are permitted to redact or remove information on their resume that identifies age, date of birth, or dates of attendance at or graduation from an educational institution.*

*We consider qualified applicants with criminal histories for employment on our team, assessing candidates in a manner consistent with the requirements of the San Francisco Fair Chance Ordinance.*

Kraken is powered by people from around the world and we celebrate all Krakenites for their diverse talents, backgrounds, contributions and unique perspectives. We hire strictly based on merit, meaning we seek out the candidates with the right abilities, knowledge, and skills considered the most suitable for the job. We encourage you to apply for roles where you don't fully meet the listed requirements, especially if you're passionate or knowledgeable about crypto!

We may ask candidates to complete job-related skills or work-style assessments as part of our hiring process. These assessments are designed to evaluate competencies relevant to the role and are applied consistently across candidates for similar positions. Assessment results are considered alongside other relevant information, such as experience and interviews, and are not the sole basis for any employment decision.

As an equal opportunity employer, we don't tolerate discrimination or harassment of any kind. Whether that's based on race, ethnicity, age, gender identity, citizenship, religion, sexual orientation, disability, pregnancy, veteran status or any other protected characteristic as outlined by federal, state or local laws.

### **Stay in the know**

[Follow us on Twitter](#)

[Learn on the Kraken Blog](#)

[Connect on LinkedIn](#)

[Candidate Privacy Notice](#)

**Apply for this Job**

Powered by **Ashby**

[Privacy Policy](#) [Security](#) [Vulnerability Disclosure](#)