

## Job Opening: DevOps / Infrastructure Engineer



Location: Berlin / Remote (within Germany) | Language Requirement: English

This is a permanent in-house (Festanstellung) position. We hire candidates already residing in Germany with a valid work permit.

### The Big Picture

Every day, our platform handles millions of customer service interactions — emails, calls, chats, voice — powered by AI that needs to respond in real time, every time. Behind every intelligent reply and every automated workflow sits infrastructure that simply cannot fail.

#### That infrastructure is your domain.

Enneo is an AI-native customer service platform. We don't bolt AI onto legacy software — the entire system was built from the ground up around AI agents, LLM orchestration, and omnichannel automation. We serve some of Germany's largest enterprises in regulated industries, process billions of tokens monthly, and are scaling fast.

None of that works without someone making sure the machines keep running, the deployments stay smooth, and the systems scale without breaking a sweat.

**Compensation:** €55,000 – €75,000 base salary depending on experience.

### Your Role

You own our infrastructure — end to end. From Kubernetes clusters and bare metal servers to observability pipelines and deployment automation: you make sure enneo is fast, reliable, and ready to scale.

This isn't a "maintain what exists" role. We're at the stage where infrastructure decisions have outsized impact. The systems you build now will carry us through the next order of magnitude of growth. You'll work closely with our engineering team and founders to shape how we run, deploy, and monitor everything.

### What You'll Do

#### Keep the Lights On — and Make Them Brighter:

You manage and evolve our production infrastructure: Kubernetes orchestration, bare metal servers for customer-facing workloads, cloud servers for internal services. You ensure uptime, performance, and security across the board.

#### Build the Deployment Pipeline:

You own CI/CD, release processes, and infrastructure-as-code. You make deployments boring — predictable, automated, and rollback-safe. When something goes wrong in production, you're the one who finds it, fixes it, and builds the system to prevent it from happening again.

#### Observability & Incident Response:

You expand and maintain our observability stack (SigNoz) so the team can debug fast and spot issues before customers do. You orchestrate LLM agents so they build dashboards, alerts, and tracing that give everyone confidence in what's running.

#### Scale the AI Engine:

Our platform routes large volumes of LLM requests (>1 bn tokens per day and growing) and relies on OpenClaw for analytics and development assistance. You ensure this layer performs reliably under load, optimize costs, and help the team iterate on AI infrastructure without friction.

### Automate Everything You Can:

We run a lot of our infrastructure management through Python-based automations and tooling. You'll build, maintain, and improve these — and increasingly leverage agentic development patterns to make our operations smarter and more self-healing. We grow fast — and with the click of a button we should be able to do so.

### Shape the Stack:

We're transparent about what we use — check out our full stack at [stackshare.io/companies/enneo](https://stackshare.io/companies/enneo). You'll have a direct voice in evolving it as we grow.

## What We're Looking For

### Must-Haves:

- Strong hands-on experience with Kubernetes — managing, troubleshooting, and scaling self-hosted clusters in production (kubeadm, k3s, RKE2, OpenShift, or similar). Cluster bootstrap, control plane, etcd, networking (CNI), storage (CSI).
- Hands-on experience with bare-metal Linux servers in production — physical servers, kernel/network/storage tuning, hardware-level troubleshooting. Solid Linux systems knowledge across both cloud and bare metal environments.
- Proficiency in Python for automation, scripting, and building internal tooling. This is not optional — Python is how we manage our infra.
- Familiarity with agentic development workflows. You understand how AI-assisted tooling and automation can be applied to infrastructure work, and you're already using it.
- Experience building and maintaining CI/CD pipelines, infrastructure-as-code, and automated deployment workflows.
- Experience managing and operating production databases — in particular MySQL at scale: replication topologies, failover (e.g., Orchestrator, MHA, ProxySQL), backup/PITR, performance tuning. Familiarity with vector databases (we use Weaviate) is a strong plus.
- Strong understanding of observability: metrics, logging, tracing using the OpenTelemetry standard. Our stack is built on SigNoz (OpenTelemetry-based).
- A true startup mindset: you're comfortable in environments where not everything is documented, things change fast, and your decisions have immediate, real impact. You don't wait for a ticket — you see what needs to happen and you do it.

### Nice-to-Haves:

- Backend Software development background
- Experience with LLM infrastructure — request routing, token management, cost optimization (e.g., LiteLLM or similar proxies).
- Experience with incident response processes, on-call rotations, and post-mortems.
- Familiarity with networking, DNS, load balancing, and CDN configuration at a production level.

## How We Work

**Impact Over Hours:** We care about what you deliver, not when you're online. Flexible working hours, remote-first within Germany, with our Berlin office available whenever you want it. We're a startup — things move fast, priorities shift, and sometimes the week has more in it than you planned. The best people here are the ones who find that energizing rather than exhausting.

**Ownership Culture:** We operate with minimal hierarchy and maximum trust. You won't find endless approval chains here. You also won't find someone handing you a playbook — we're still writing it. We're

## Job Opening: DevOps / Infrastructure Engineer



looking for people who take a problem, own it, and solve it — and who are energized by that kind of freedom.

**Close to the Founders:** This is not a role where you disappear into a department. You work directly with our founding team and have a real voice in shaping how enneo grows. Your infrastructure decisions directly influence our product capabilities. That kind of influence comes with real responsibility — and we wouldn't have it any other way.

### Why This Role Matters

AI platforms live or die on their infrastructure. The smartest model in the world is worthless if it can't respond in time, can't scale under load, or goes down at 2am on a Tuesday. You're the person who makes sure that never happens — and who builds the systems so it couldn't happen even if everyone else is asleep.

If you want to own the infrastructure layer of a fast-growing AI company — not maintain someone else's setup, but truly architect and scale it — this is your role.

**Ready?** Send your application to [kariere@enneo.ai](mailto:kariere@enneo.ai). Let's build the foundation together.

enneo is an equal opportunity employer. We welcome applications from all qualified candidates regardless of gender, nationality, ethnic origin, religion, disability, age, or sexual orientation.