



How Vet-AI reduced annual DevOps costs by £80k+ and increased velocity by 20x



VET-AI
AI
FOUNDED 2018

From Kubernetes headaches to high-velocity shipping and cost savings

04/14/25

4 MIN READ




MARCUS KOHLBERG



with clinically validated
technology.

Book a demo

In partnership with



Building the future of pet health

Vet-AI is on a mission to empower pet owners with smarter, clinically validated technology for managing their animals' health. With a web app and native mobile apps for both iOS and Android, the team is creating a unified platform that combines AI-driven diagnostics, continuous care, and seamless user experience.

But when it came time to replatform their backend infrastructure, Vet-AI faced a familiar challenge: scaling their team quickly without being dragged down by DevOps complexity.

Quantifying the impact

Encore helped Vet-AI dramatically reduce infrastructure overhead and developer distractions.

- **£80k+ annual reduction in DevOps cost** – replacing the need for a full-time DevOps engineer, a ~£8k/month cost, with a £1k/month platform
- **20x release velocity** – Over 200 deployments/month, up from 1–2 per week
- **Onboarding 20 engineers in a single day** – with fully working local environments and zero friction

- **Faster time-to-market** – smaller, more frequent releases improved feature velocity and reduced risk
- **Zero infra firefighting** – no more YAML pipelines or debugging staging environments

From Kubernetes headaches to high-velocity shipping

Before Encore, Vet-AI's stack was based on Kubernetes on Google Cloud Platform. It worked, but it came with the kind of DevOps tax that slows startups down. Rather than keep investing in an increasingly complex setup, the team chose to replatform entirely.

We looked at things like Cloud Run, App Engine, and of course managing Kubernetes ourselves. But the effort required to spin up environments on every PR, wire up all the CI/CD pipelines, and maintain infrastructure didn't make sense. With Encore, we could just focus on building and still keep using Google Cloud Platform.

Rob Gray, Head of Engineering at Vet-AI

That focus paid off.

We went from shipping once or twice a week to over 200 deployments a month. That's a 20x increase in release velocity.

Charles Winer, Senior Engineer at Vet-AI

Encore's out-of-the-box features like dynamic environments, automatic tracing, and self-documenting APIs transformed the team's development workflow. New engineers could clone the repo, run a single command, and be up and running with a fully working API and web environment in minutes.

We set up a one-day onboarding for 20 external contractors, and the engineers

could just clone the code, install Encore, run it, and the API's live. Super smooth and low-stress. The team said it was one of the fastest onboardings they'd seen.

Rob Gray, Head of Engineering at Vet-AI

Fewer distractions, more product

By eliminating DevOps as a bottleneck, Vet-AI turned every engineer into a product-focused contributor.

Before, our backend engineers were also our infra engineers. It took time and focus away from what mattered. Now we don't even need a DevOps team. Encore handles secrets, custom domains, and cloud integrations. It's just all there.

Rob Gray, Head of Engineering at Vet-AI

For a team with mobile apps, the visibility across services made a huge difference too.

Our mobile developers love having access to Encore Cloud's dashboard. They can inspect traces, see API changes in real-time, and understand how things fit together, without us needing to manually maintain documentation or telemetry.

Charles Winer, Senior Engineer at Vet-AI

Built for scale, ready for global growth

Looking ahead, Vet-AI is expanding internationally and Encore is already helping them lay the groundwork.

One of the reasons we chose Encore was the ability to spin up environments in new

regions quickly. That's huge for our global ambitions. The fact that you can just add a new environment, pick your cloud platform, pick your region, and it'll just spin it up and then all your deployments to main will just deploy into each region automatically for you. It just takes a lot of the headache out of going multi-region

Rob Gray, Head of Engineering at Vet-AI

Takeaways

Vet-AI's experience offers a clear takeaway for startups:

If you're rebuilding or even just getting started, Encore lets you skip the pain of DevOps and get straight to building. You'll ship faster, onboard easier, and focus more of your time on your actual product.

Rob Gray, Head of Engineering at Vet-AI

You can absolutely do all of this yourself, but why would you want to?

Charles Winer, Senior Engineer at Vet-AI

Results at a glance

- £80k+ annual DevOps cost reduction
- 20x release velocity
- 20 engineers onboarded in a day, fully productive from day one
- Zero DevOps bottlenecks, total focus on product
- Global-ready infrastructure, no manual setup required

Encore enabled Vet-AI to replatform with speed, scale their team overnight, and ship AI-powered features faster than ever. All without writing a single YAML file.