

Production ML. Built to last.

Clavenix helps healthcare and fintech teams move Deep Learning models from notebooks into dependable production systems on GCP and AWS. Engagements are fixed-scope, milestone-based, and designed for your team to own the resulting infrastructure.

THREE WAYS TO ENGAGE

MLOps Readiness Audit

01 / FIRST STEP

A 3–5 day review of model artefacts, infrastructure, team workflows, compliance gaps, and deployment blockers. You receive a written roadmap with no obligation to continue.

Typical output: prioritised findings and deployment roadmap

Core Deployment Project

02 / MOST COMMON

A production deployment for one or two models on GCP or AWS, including containerisation, CI/CD, drift monitoring, HIPAA-aware inference architecture, documentation, and 30-day post-deployment support.

Typical timeline: 4–8 weeks

Complex Multi-Model Project

03 / ENTERPRISE SCALE

Shared infrastructure for multiple models and complex pipelines, with orchestration, automated retraining triggers, advanced monitoring, HIPAA- or SOC2-grade controls, and team handover.

Timeline: scoped to complexity

FROM UNCERTAINTY TO OWNERSHIP

Clear scope. Practical delivery.

The right engagement starts with an honest assessment of where your deployment stands. Every project is scoped to your models, cloud environment, compliance needs, and operating team.

01

Understand

Review the model, current infrastructure, team workflows, and constraints.

02

Scope

Define deliverables and milestones before deployment work begins.

03

Deploy

Build reproducible pipelines, inference infrastructure, and monitoring.

04

Handover

Document the system and equip your team to operate it independently.

Not sure where to start?

Begin with a Readiness Audit. Its written findings help clarify whether your next step is a focused deployment, a multi-model programme, or internal remediation first.

Discuss your project: hello@clavenix.com