




// WHY HELMCODE

Add open models to your AI stack.

Without adding another
infrastructure project.

Use production-ready open models alongside
your existing frontier providers, through a single
inference platform.

YOUR AI STACK



an additional inference option · not a migration

// WHAT IS HELMCODE

Open model inference, operated for you.

MODELS


Access leading open models through one API.

INFRASTRUCTURE

We operate the inference infrastructure behind them.

ENTERPRISE

Deployment, security and support for production workloads.



Not every workload needs a frontier model.

CHOICE

Use the model that fits the workload instead of forcing every workload through the same provider.

CONTROL

Choose how and where inference runs as requirements change.

ECONOMICS

Match model capability and infrastructure cost to the actual task.

Open models don't need to replace your current AI stack. They give it more options.

further reading → [AI Sovereignty: who controls your company's intelligence?](#)

// CLIENTS

Clients working with us




AI engineering consultancy



AI startup




public sector



travel eSIMs · 1,000 developers



AI for accounting · €100K/month → €7.7K/month



enterprise open source software

// ENTERPRISE READY

Enterprise ready.

Full security, privacy and compliance documentation available during evaluation.

DATA



Zero inference-content logs
No training on customer data
EU processing for Helmlcode-operated open models

SECURITY



TLS 1.3
API key controls
Dedicated / isolated deployments

COMPLIANCE



GDPR
DPA available
Subprocessor information

DEPLOYMENT




Managed EU infrastructure
Dedicated
On-premise / air-gapped

// PORTABILITY

Switch providers, not your stack.

Helmcode operates open models. It doesn't own them.



Your inference provider should be a decision you can change.

note: providers differ in quantisation, context limits, latency
and supported features

// START WITH A POC


Start with one workload.

01	02	03	04	05
Pick a workload	Establish the current baseline	Test open models with Helmcode	Compare quality, latency, cost and security	Decide with data

No platform migration required.

// FROM POC TO PRODUCTION

A low-risk path to production.



Start with one workload. Expand only where Helmcode proves useful.



Start with one workload.

enterprise evaluation · pedro@helmcode.com