

Build with OpenAI on AWS – from Day One



Start your AI & agentic project on AWS the right way: the latest OpenAI models, an OpenAI launch partner and generous AWS funding.

With OpenAI now GA on Amazon Bedrock, you can start building directly on AWS. Same AWS account, same data, same security model, same procurement, with the latest OpenAI models available on day one, alongside the latest AI tools and solutions from AWS.

As an OpenAI launch partner with 500+ GenAI and Agentic launches since 2023, Loka delivers a six-week build engagement that takes you from concept to a functioning pilot, with generous AWS funding for qualified customers.

9.7 out of 10
(80+ AWS Customer Reviews)

NOW LIVE: OpenAI models GA on Amazon Bedrock. Loka is a global launch partner.

Why Build with Loka on AWS

OpenAI on day one

The latest GPT models available in your AWS account from the start. No separate vendor relationship, data sprawl, or separate billing.

Build with the experts

Loka brings 500+ GenAI and Agentic launches and AWS Partner of the Year credibility to your first production deployment.

One cloud, one stack

Your data, vector DB, agents, identity, observability, and deployment all live in the same place from project zero.

Model optionality

Benchmark OpenAI, Anthropic Claude, Amazon Nova, and open source per workload — pick what's right per use case, swap freely.

How It Works

- **Phase 1: Use Case Shaping | 1-2 weeks**
Define the right first use case, success criteria and evaluation framework. Pick the right model candidates.
- **Phase 2: Architecture + Pilot | 3-4 weeks**
Build a functioning pilot on Bedrock with OpenAI and at least one comparison model. Live evals against real data.
- **Phase 3: Implementation Roadmap | 1 week**
Identity, networking, observability, cost controls, agent loops, deployment automation and SLO definition — the plan to take your pilot to production.
- **Next step: Production**
If the pilot fits, roll into production with Loka, often with AWS funding.

What You Get

- Functioning pilot + initial architecture**
A working model and initial AWS architecture running the latest OpenAI models.
- Eval framework + benchmark data**
Reproducible evals across the model candidates you tested, with real data. Repeat as models evolve.
- Implementation roadmap**
Identity, networking, observability, cost controls, and SLOs sequenced to take your pilot to production.

What Customers Say

“Working with Loka was an easy yes for Cameo. Our partnership allows us to move faster and punch above our weight class with more complex LLM implementations—the kind that remove friction and keep our customers coming back for more”

Dom Scardinaro
CTO at Cameo

Who Is This For?

- AWS customers planning their first GenAI initiative
- Teams with executive mandate for GenAI but no clear starting point
- Product teams who want OpenAI models without managing OpenAI as a vendor
- Engineering leaders consolidating GenAI on AWS from day one to avoid future migration
- Teams with employees who use OpenAI with little-to-no organizational oversight

[Click to book your OpenAI Assessment](#) 