



OpenAI on AWS Migration Assessment

Know what can move now and what should wait. Get a one-week assessment complete with feature gap triage, 100% funded by AWS.

Bring your OpenAI workloads to AWS without giving up the latest models your teams already use. Before you migrate anything, feature gap triage tells you in plain terms what ports cleanly, what needs a workaround and what should wait, including how your data isolation and residency posture changes on Bedrock. Bedrock supports a subset of OpenAI capabilities at launch. As one of a select few OpenAI-on-Bedrock launch partners, Loka reviews your production workloads against the current Bedrock feature set.

From there we benchmark your setup on Bedrock, map your architecture from API surface to model routing and data flows, and build a roadmap that won't disrupt your team. You get it in one week. If the plan looks good, Loka can own the full migration, done in under five weeks for most projects.



9.7 out of 10
(80+ AWS Customer Reviews)

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

Why Assess with Loka

Deep feature gap triage

Inventory your workloads before signing off on any migration budget. No guesswork, no scope creep on day one.

Decide workload by workload

Every blocking gap surfaced with a go, wait, or re-architect recommendation, and an AWS-native alternative for each feature.

Unify your stack

One bill, one IAM, one audit trail. The latest models, vector DBs, agents, observability, and data in one place, with AWS funding available for qualifying workloads.

Backed by production, not pilots

Half of our 500+ GenAI and agentic projects since 2023 are in or moving to production, most for customers in regulated industries.

How It Works

Phase 1: Assessment & Workload Inventory | 2-3 days

Inventory your production OpenAI workloads: models, API surface, features in use, data flows, and any regulated or compliance requirements.

Phase 2: Gap Mapping | 3-4 days

Map each workload against the current Bedrock feature availability table. Surface any blocking gaps, for example hosted web search, file search, computer use, remote MCP servers, WebSocket connections, or audio input, and identify AWS-native or alternative options for each.

Phase 3: Recommendation | 1-2 days

A per-workload go, wait, or re-architect recommendation, delivered in a plain readout deck your team can act on immediately.

What You Get

Benchmark vs your OpenAI workloads

Cost, performance, security, and latency on Bedrock measured against your existing OpenAI setup.

Blocking gap list with alternatives

Every gap identified with a clear AWS-native or workaround option, or a plain wait recommendation per feature.

Per-workload recommendation roadmap

Go, wait, or re-architect for each workload, with the reasoning your team can share with leadership and audit later.

Executive go/no-go and option to migrate

A board-ready summary with the option to move fast into full migration with Loka.

Who Is This For?

- Teams unsure whether their OpenAI workloads can migrate to Bedrock before committing budget
- Companies running OpenAI in production and weighing whether to migrate before committing resources
- Teams who have heard about Bedrock feature gaps and want to know exactly which ones apply to them
- Teams whose InfoSec or procurement org has flagged OpenAI direct
- Teams that want their models in the same place as the rest of their apps and data, under one bill and one audit trail

What Customers Say

"Loka helped us do two things at once: move a lot of our machine learning workflows onto AWS very quickly and modernize our DevOps stack substantially."

Catalin Voss

CEO, Ello, #1 AI Reading App for Children

[Click to book your rapid OpenAI on AWS Assessment](#)