

OpenAI Migration Fast Track

Your OpenAI Workloads, Running on AWS. Assessed in a Week.
Migrated in Five. 100% Funded by AWS.



For companies running OpenAI in production, the migration question has always come with friction: how long, how much, how disruptive. The OpenAI Fast Track to AWS answers all three.

Keep the OpenAI models your team already trusts and move the control plane to AWS. Loka swaps your application to the Bedrock OpenAI client, points it at the right US Region and model, and validates that behavior matches your current baseline. You get AWS-native billing, IAM, audit logging, and KMS encryption, with no change to the model you rely on.

As a global OpenAI-on-Bedrock launch partner with 500+ GenAI and agentic projects since 2023, Loka runs both phases end to end. One engagement, one team. AWS security, one bill, the same models, live in five to six weeks.



9.7 out of 10
(80+ AWS Customer Reviews)

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

Why Fast Track with Loka

\$0 Cost for Qualified Customers

AWS covers your full assessment cost and up to 100% of your migration for qualified customers. Qualification is confirmed during your discovery call.

AWS governance inherited

IAM, KMS encryption, CloudTrail logging, and VPC isolation applied automatically from day one.

Two Entry Points

OpenAI direct or Azure OpenAI. One fast track to move you to production the right way.

How It Works

Phase 1: Fast Track Assessment | Up to 1 week

Inventory your OpenAI workloads, API usage, prompts, models and data flows. Map the target AWS architecture. Confirm qualification and deliver a migration plan ready to execute.

Phase 2: Fast Track Migration | 4 to 5 weeks

Full migration of qualified OpenAI workloads to Amazon Bedrock. Swap to the BedrockOpenAI client with correct Region and model configuration. Run behavior-parity validation against your OpenAI baseline. Verify billing, IAM, KMS and CloudTrail end to end.

Qualification Criteria

- Minimum \$5,000 per month in OpenAI API spend, or at least one active production OpenAI workload
- Workloads are actively running in production, not pilots or proofs of concept
- Migration scope is OpenAI workloads only. Other workloads are estimated and funded separately
- Company meets AWS funding eligibility requirements, confirmed on the discovery call

What You Get

Fast Track Assessment report

Workload inventory, AWS architecture map, compliance posture, and a migration plan delivered in one week.

Full production migration

Qualified OpenAI workloads migrated to Amazon Bedrock, tested, and live on AWS. Correct Region and model configuration, tested and validated against your current baseline.

Post-migration handoff

Runbooks, observability setup, and knowledge transfer to your engineering team.

Additional funding surfaced

If other workloads are in scope or additional AWS programs apply, Loka identifies and unlocks those separately.

What Customers Say

"We had a very aggressive timeline. Building on AWS with Loka allowed us to scale and get to market very quickly. We are now live on about 2,000 providers and they are seeing a reduction in times per visit between 5 to 15 minutes per encounter."

Chris Hoffman

Chief Architect at Experity Health

Who Is This For?

- Teams whose InfoSec or procurement org has flagged OpenAI direct API usage
- OpenAI in production - AWS billing and governance, zero code changes.
- Any US-based OpenAI user who wants to unify AI spend inside AWS
- Engineering leads who need IAM, audit logging and US data residency applied to AI workloads

[Click to book your OpenAI Fast Track Migration](#) 