EDT & PARTNERS | amber

From Internal Tool to Commercial Platform: An Agentic AI Transformation

Transforming Education with Purpose

Introduction

Amber is an AWS Advanced Tier Services Partner operating across Peru and the United States. The company builds cloud-native applications and AI solutions for financial services organizations, startups, and government clients, and holds the AWS AI Services Competency, a distinction held by only a small number of partners in the Latin American region.

By late 2025, Amber had developed **TARS: a proprietary platform designed to automate cloud operations, incident management, and internal workflows using a coordinated network of AI agents.** The ambition was clear. Evolve TARS from an internal tool into a commercial product available on AWS Marketplace, and position Amber as the leading agentic AI partner in Latin America and beyond. The challenge was not vision, it was architecture.

EDT&Partners delivered the Agentic AI Module of the AWS Partner Transformation Program and within four months of completing it, Amber had validated TARS's production architecture, launched a new commercial consulting product, and built a pipeline of over $150,000 in agentic AI opportunities.

Results at a glance: $150K+ AI pipeline · 5+ qualified opportunities in four months · 60% faster concept-to-prototype · ~5x development speed

Challenge: Strong Foundations, an Even Bigger Ambition

The Amber team already had hands-on experience with the core building blocks: Amazon Bedrock for running AI models, vector databases on AWS for storing and retrieving knowledge, and the Strands Agents SDK for building individual AI agents. Moving from capable individual agents to a production-grade, commercially scalable multi-agent platform is a different challenge entirely.

Getting TARS to market meant resolving the following:

Depth in Amazon Bedrock AgentCore. AgentCore is the AWS service suite that handles the critical infrastructure for running AI agents at scale: secure deployment, tool connectivity, identity management, and operational monitoring. Amber had the ambition to deploy at scale. Deepening mastery across the full AgentCore portfolio was the next step.

Multi-agent orchestration. TARS is not a single agent. It is a system where a central coordinating agent delegates tasks to a set of specialized sub-agents, each handling a different function. Designing that coordination layer reliably, with proper fallback logic and supervision, needed to be validated before going to production.

Long-term memory. For TARS to learn from operational history and maintain context across client sessions, a validated approach to persistent memory was needed. How long should information be retained? How do you handle data privacy across clients? How do you design the memory architecture so agents can use what they know? These were open questions.

Commercial readiness. Before TARS could be listed on AWS Marketplace and sold to multiple clients simultaneously, the architecture needed to meet enterprise-grade standards: data isolation between clients, robust security controls, and the ability to scale without rebuilding from the ground up.

Resolving these was the bridge between a strong internal product and a market-ready platform.

Solution: Months of Architectural Decisions, Resolved in Weeks

EDT&Partners delivered the Agentic AI Module as four structured sessions covering the full agent development lifecycle on AWS. The program was not a generic curriculum. It was designed around Amber's specific priorities and applied directly to the architecture they were building.

The working case throughout was TARS itself. Named after the autonomous robot in the film Interstellar, TARS embodies the same principle: an intelligent system that works alongside human teams to handle complex tasks autonomously while remaining under human oversight. Using TARS as the live reference meant every concept was tested against something real.

By the end of the program, Amber had a validated, production-ready architecture built around four core capabilities.

The orchestration layer. A central coordinator agent delegates tasks to a set of specialized sub-agents, each focused on a specific function: monitoring infrastructure, diagnosing problems, executing remediation, and handling documentation. The delegation logic, fallback behavior, and supervision structure were all confirmed.

The memory architecture. TARS uses a three-layer approach to contextual memory: short and medium-term session memory handled by AgentCore Memory, and long-term organizational knowledge stored and retrieved via Amazon S3 Vectors and Bedrock Knowledge Bases. This was one of the most significant outputs of the program.

The incident lifecycle. TARS can now manage the full journey of a cloud support incident autonomously: detecting the alert, creating and classifying a ticket, diagnosing the issue, recommending or executing a fix, communicating with the team, and generating post-incident documentation. Human sign-off is required only at the final validation step.

The integration layer. TARS connects to the tools Amber's clients already use: Monday.com for ticket management, Slack for team communications, AWS CloudWatch for infrastructure monitoring, and internal systems via a secure tool gateway. Agents interact with these systems in live production environments.

"The architecture patterns we validated during the program, particularly the three-layer memory system and the event-driven multi-agent orchestration, became the foundation of TARS. Without that structured exploration, we would have spent months iterating on architectural decisions that the program helped us resolve in weeks. The combination of AgentCore Runtime for secure deployment and Strands SDK for agent logic gave us a stack we can confidently take to production and, eventually, to AWS Marketplace."

— Victor Cueva, VP of Engineering, Amber

Business Impact: From Technical Validation to Commercial Momentum

Metrics at a glance

- Qualified agentic AI opportunities | **5+ within 4 months**
- AI-specific pipeline | **$150K+** identified
- New commercial product | Amber Ignite launched
- TARS design-to-production timeline | Estimated **60% reduction**
- Development velocity | Approximately **5x improvement**
- Fully automated post-incident reporting | Projected **30% faster incident resolution** times

The impact went further than architecture. The program gave Amber the confidence and the framework to productize what they had learned and to take a new position in the market.

Amber Ignite
Amber Ignite is a direct product of what was built and validated during the program. It is a structured four-week consulting engagement, that transfers AI development competencies and agentic AI capabilities to enterprise clients in Peru, LATAM, and beyond. The methodology combines modern software development practices with hands-on agentic AI workshops covering the same stack validated in the program. Amber Ignite is being prepared for listing on AWS Marketplace as a Professional Services offering.

Pipeline and opportunities
Prior to the program, Amber had no structured pipeline in agentic AI. Within four months of completing it, five or more qualified opportunities had been identified, totaling over $150,000. Amber Ignite functions as a pipeline engine in its own right: workshop clients convert into architecture consulting and implementation projects.

TARS development
The program is estimated to have reduced TARS's design-to-production timeline by 60%, by validating critical architectural decisions early rather than iterating under pressure. Projected operational improvements for Phase 1 include a 30% reduction in average incident resolution time and fully automated post-incident documentation.

Development velocity
Internal teams adopted Spec-Driven Development with Kiro during the program, a structured approach to software planning and code generation that produced approximately a 5x improvement in development cycle speed. The Amber Ignite delivery model is designed for delivery by trained consultants, meaning the methodology scales beyond the founding team.

"The program gave us the structured framework we needed to accelerate TARS from concept to a validated, production-ready architecture. We already had experience with Bedrock, S3 Vectors, and Strands, but what we lacked was depth in AgentCore, particularly in multi-agent orchestration and medium to long-term context memory. The program closed that gap. It also gave us the confidence to productize that knowledge into Amber Ignite, our own consulting methodology. We went from being consumers of GenAI services to being builders and educators. That shift is what defines Amber's position today: we don't just implement AI for our clients, we transfer the capability so they can build autonomously. The program was the catalyst."

— Axel Echevarría Piérola, CEO, Amber

The Shift That Defines Where Amber Stands Today

Amber came into the program with strong foundations and a clear commercial ambition. The AWS Partner Transformation Program, delivered by EDT&Partners, provided the structured framework to turn that ambition into a validated, production-ready platform and a new market position.

TARS has a confirmed roadmap toward AWS Marketplace. Amber Ignite is generating pipeline and transferring capability to clients across the region. And Amber has established knowledge democratization as a core philosophy, one that sets it apart from the growing number of partners entering the agentic AI space and will define their market position for years to come.

*AWS partners looking to build agentic AI capability can reach out to Jasmine Thakkar, AWS PTP Global Lead (**jazzth@amazon.com**), or Alessandro Bilotta, Director of Program Management at EDT&Partners (**abilotta@edtpartners.com**).*

EDT&
PARTNERS

www.edtpartners.com
+44 1274089806 (Europe)
info@edtpartners.com