CASE STUDY

Industrializing AI for Pharma R&D

How a leading global pharmaceutical company replaced a fragmented stack of AI tools with one governed, self-service platform — automating the model lifecycle, scaling across cloud and on-prem, and giving every scientist a role-based multi-agent workspace for enterprise knowledge.

THE CLIENT

Our client is a leading global pharmaceutical company engaged in the research, development, manufacturing and delivery of innovative medicines. The organization operates across North America, Europe, Asia-Pacific and emerging markets.
Its teams span the entire pharmaceutical value chain — drug discovery, clinical development, regulatory approval, manufacturing and commercialization — and combine deep scientific expertise with real-world healthcare insight to deliver sustainable value to patients and health systems worldwide.

THE CHALLENGE

As the client's AI and platform initiatives scaled, two pressures compounded each other.At the model layer, building, training, and operating production-grade AI/ML at global R&D scale required scarce data-science talent, repeated tooling effort, and rising cloud and operational costs. At the user layer, information, models, applications, and workflows were fragmented across systems and tools, making it hard for business and technical users to get role-relevant answers or launch new GenAI capabilities without weeks of approvals.

Leadership wanted a single, governed home for AI: one platform that could automate the AI/ML lifecycle, deploy intelligently across cloud and on-prem, and give every persona, a self-service multi-agent workspace tuned to their role and intent.

OUR APPROACH

We partnered with the client's R&D and platform teams to design and build AI on AMP — a unified, self-service AI platform that replaces fragmented point tools with a single, governed layer for the full AI lifecycle.

The goal was deliberately broader than another GenAI tool. AI on AMP had to industrialize how AI gets built, deployed and consumed across the enterprise — bringing AutoML, hybrid scaling, agentic orchestration and role-aware information retrieval together in one place, without forcing the R&D community to abandon the tools they already trusted.

INSIDE THE SOLUTION

1. A Governed Sandbox for AI/ML Experimentation
2. Industrialized AI/ML Automation
3. Open-Source Model Onboarding, Without Lock-In
4. Hybrid Deployment — Cloud and On-Prem, by Design
5. An "AI for AI" Optimization Layer
6. A Role-Based Multi-Agent GenAI Workspace
7. Client-Specific Adaptors for Industry Tools

OUTCOMES

AI on AMP moved the client from a fragmented toolset and manual processes to a centralized, self-service AI ecosystem. Early outcomes include:

- A single, governed home for AI/ML and GenAI
- Material efficiency gains on AI workloads
- Faster, role-relevant insights for business users
- Streamlined deployment and coding assistance for developers
- Confident, self-service experimentation
- A scalable, governed foundation

WHY IT WORKED

Three design choices defined the outcome. First, we treated this as a platform problem, not a model problem. Second, we kept the architecture model-agnostic and deployment-agnostic, so the client could absorb open-source innovation at its own pace, across both cloud and on-prem, without being locked to a vendor. Third, we built governance, role-based routing and an AI-for-AI optimization layer in from day one