



Xaltris Technologies

Powered by AI, shaped by OI

*Original Intelligence

Case Study

AI Assisted Governed Multi-Cloud Lab Generation Platform

Overview

This solution is a governed, agentic AI platform that plans and orchestrates the creation of production-ready multi-cloud labs from natural-language learning objectives, with humans retained in the loop for review and execution. Grounded in approved service catalogues via a RAG architecture, it replaces ad-hoc lab authoring with a controlled, auditable, end-to-end workflow.

The Problem

Hands-on cloud training is difficult to scale. Lab creation is often manual and inconsistent, governance and clean-up are unreliable, and there is no durable system of record for what was generated or deployed. These gaps lead to fragile learning experiences, elevated security and cost risk, and unnecessary operational overhead.

The Solution

The platform introduces a governed, agentic workflow that transforms learning objectives into structured, reviewable lab packages. Labs integrate bi-directionally with existing LMS platforms - ingesting course context and learner state, while pushing back completion data, artifacts, and outcomes. All outputs are grounded in approved multi-cloud services, refined through human-in-the-loop review, and tracked across their full lifecycle - enabling quality, traceability, and confidence at scale.

Key Capabilities

Structured Lab Generation

From a short natural-language objective, the system generates a complete lab package, including a structured lab specification, step-by-step learner instructions, a weighted grading rubric, and infrastructure provisioning artifacts. All outputs are execution-ready, human-reviewable, and designed for instructional clarity.

RAG-Grounded Multi-Cloud Intelligence

A curated AWS, Azure, and GCP service catalogue is indexed in a vector database. During generation, only approved services, patterns, and reusable snippets are retrieved and injected into the model context - ensuring alignment with organizational standards and known-good practices.

LMS Integration and Learning Flow Synchronization

The platform integrates bi-directionally with Learning Management Systems (LMS), ingesting course structure, learner context, and assignment metadata, while exporting lab availability, completion status, & assessment results. This allows labs to function as native learning units within existing curricula, without duplicating workflows or fragmenting learner data.

Governed Execution and Lifecycle Control

Before provisioning, users can review and refine generated outputs. Automated guardrails enforce tagging standards, service constraints, and permission boundaries. Provisioning and clean-up follow a controlled orchestration flow designed for safety, repeatability, and cost control.

Auditability and System of Record

Every provisioning action is recorded - including what was created, when it was created, who initiated it, and which lab it belongs to. This establishes a durable system of record for cost attribution, compliance reporting, and continuous improvement.

Technology Stack

- **Frontend** : React + Vite + Tailwind
 - **Backend** : Python (FastAPI)
 - **LLM** : Google Gemini
 - **RAG / Vector Database** : ChromaDB
 - **System of Record** : PostgreSQL
-

Outcome & Value

The platform enables fast, reliable, and governed lab generation at scale - reducing authoring effort, improving consistency, and embedding safety and traceability by design. With full lifecycle visibility and a durable system of record, it provides a strong foundation for confident execution today and analytics-driven improvement over time. The solution is well suited for organizations delivering hands-on cloud training across teams, clouds, and cohorts.