



# A Case Study on Enabling Enterprise- Scale AI with Unified Analytics Infrastructure

Service Provider: Whiteklay



## CLIENT OVERVIEW

A global enterprise with diverse AI initiatives across departments sought a unified solution to streamline their machine learning lifecycle. From data preparation to model tuning and GenAI deployment, the goal was to create an end-to-end AI environment capable of operating at enterprise and supercomputing scale.

## OBJECTIVE

To build a unified, enterprise-grade AI infrastructure that simplifies the end-to-end machine learning lifecycle — from data ingestion to LLM deployment — while ensuring scalability, governance, and cross-environment orchestration across cloud, edge, and on-prem systems.

## PROBLEM IDENTIFICATION



Disconnected tools for data management, model training, and inference



Difficulty scaling LLMs due to GPU and deployment constraints



Inconsistent data lineage, versioning, and orchestration



Hybrid infrastructure complexities across edge, cloud, and on-prem



## SOLUTION APPROACH

1. Deployed an AI/ML lifecycle platform with modular and scalable components
2. Integrated a unified data fabric and federated SQL layer for consistent access
3. Enabled GenAI workflows with optimized models and pipelines
4. Used GPU-powered infrastructure to accelerate LLM training and inference
5. Orchestrated end-to-end workflows with Kubeflow and enterprise controls

## BUSINESS OUTCOMES

Faster AI deployment across hybrid environments

Quicker LLM training with GPU-optimized stacks

Improved model accuracy with real-time data access

Stronger AI governance with full lifecycle visibility

Scalable architecture for GenAI across teams

**Schedule free demo:**  
[Contact Us – Whiteklay](#)