



CASE STUDY

# Transforming Inspire Clean Energy: Streamlined Data Processing and ML Ops with Spark

## Executive Summary

*Inspire Clean Energy, a US retail energy company specializing in renewable wind, solar, and hydro power, sought assistance from New Math Data (NMD) to enhance their development process and scale their workloads. NMD recommended, designed, and implemented a Spark processing platform using Databricks, integrated with Inspire's existing resources and AWS accounts. This integration included migrating Inspire's existing MLflow server to Databricks, leveraging its managed MLflow functionality for a unified experience in model development, experiment tracking, management, and deployment. The platform was securely integrated with Inspire's SSO system, providing streamlined access for their data scientists and engineers. Additionally, NMD improved Inspire's MLOps capability, ensuring reliability, reproducibility, and portability of ML models across their development environments. This strategic enhancement included configuring internal networking to accommodate both Databricks and existing AWS services like RDS within the same VPC, optimizing resource utilization and security.*

## Customer Description


Inspire Clean Energy is a US retail energy company, specializing in renewable wind, solar and hydro power for its clients, while providing a predictable pricing model.



## Description of Service

Inspire currently has a sophisticated data science and engineering team and has built a data platform comprising Snowflake, dbt pipelines, Airflow, Kubernetes.

Inspire wanted to migrate their existing MLflow server to Databricks, which has built-in, managed MLflow functionality. Databricks offered a one-stop experience for model development, experiment tracking, model management, and deployment. MLflow utilized S3 storage for its model artifacts.



NMD recommended, designed, and implemented a Spark processing platform utilizing Databricks and integrating it with their existing resources and AWS accounts. The platform was also integrated to Inspire's existing SSO system to provide secure and streamlined access for their team of data scientists and engineers.

## Description of Solution

Databricks is compatible with multiple cloud platforms, including AWS. The first step in implementing Databricks on AWS is setting up the internal networking on Databricks. Databricks can either be deployed in its own, out-of-the-box VPC which includes all the necessary subnets, route tables, and security groups, or it can be configured with customized networking. Since Inspire already had an RDS instance deployed in their AWS environment, they wanted Databricks to be deployed into the same VPC as the RDS instance, so we chose the customized networking option.

Inspire's developers needed a service to develop their models. Databricks' notebooks provided a space for them to do so. In order to run a Databricks' notebook, a Databricks cluster must be running. Databricks' compute (clusters) is integrated with AWS and uses EC2 instances to power its clusters. In order to run a Databricks' cluster, utilizing AWS EC2 instances, an IAM instance-profile role must be configured.

NMD designed and implemented/integrated a Spark processing environment based on the Databricks platform into their existing data platform and operations, which included – infrastructure definition using Terraform, integrated with Snowflake, GitHub, and Airflow running on Kubernetes in existing AWS accounts.

NMD also designed and implemented a robust MLOps pipeline that enables Inspire to have reproducible and scalable model training environments that streamline their model training, validation, and promotion processes.

To summarize, the AWS services utilized at Inspire were S3 for storage, VPC(including subnets, security groups, routing tables, etc.), EC2 for compute, and IAM for role-based access to various AWS services.

### Description of Outcome

With the successful deployment of the Databricks infrastructure, Inspire Clean Energy has seen dramatic improvements in dataset scalability and a more than 50% reduction in model training times. The integration has enhanced Inspire's ability to manage model lifecycles using Databricks' managed MLflow, with S3 utilized for storing model artifacts. Customized networking solutions were implemented to integrate Databricks with Inspire's RDS in AWS, utilizing existing security groups and subnets to ensure optimal performance and security. This strategic use of AWS capabilities has streamlined the model development and deployment process, significantly enhancing productivity and security.