

nClouds | AWS Case Studies

Global Energy & Petrochemical Leader

How nClouds helped a global energy and petrochemical leader accelerate DevOps and enhance its machine learning platform's scalability to handle petabytes of data.

Industry

Renewable Energy

Location

United States

Challenge

Accelerate DevOps and enhance the scalability of its machine learning platform to handle petabytes of data.

Featured Services

DevOps Services,
Containerization, Migration
Services, High Availability

About Global Energy & Petrochemical Leader

The company is a United States subsidiary of a global energy and petrochemical leader focused on renewable energy.

Benefits Summary



Accelerated DevOps



Enhanced scalability

Challenge: Accelerate DevOps and enhance scalability for its next-generation machine learning platform to handle petabytes of data.

In the company's traditional workflows, the Global Energy & Petrochemical Leader's data was locked in application silos. The company was launching a new machine learning (ML) platform for data analytics for use by its customers and trading partners globally. Due to tight deadlines, they needed expert help to accelerate their DevOps initiatives. And, they wanted to ensure that AWS best practices were applied to scale the environment to handle petabytes of data.

Why AWS and nClouds

The company's head of product development was impressed with the work that nClouds, an AWS Premier Consulting Partner, had done for him previously. Based on nClouds' AWS technical expertise in DevOps, containerization, and migration, he asked nClouds to help build and deploy their AWS infrastructure to meet the company's aggressive deadlines. nClouds served as a trusted partner of the organization's distributed team across five countries.



nClouds has been the trusted partner for my distributed team, across five countries. Within our first engagement, they partnered with us on a critical project that enabled us to automate and deploy services from our next-generation energy platform regionally and globally. We could not have hit these targets if not for their well-architected and well-executed work. Exceptionally talented folks!"

— CTO,
**Energy Platform, Global
Energy & Petrochemical
Leader**

Global Energy & Petrochemical Leader is leveraging several Amazon Web Services:

- **Amazon CloudWatch** - Monitors applications, responds to system-wide performance changes, optimizes resource utilization, and presents a unified view of operational health.
- **Amazon EFS Container Storage Interface (CSI) Driver** - Provides a CSI interface that allows Amazon EKS clusters to manage Amazon EFS file systems' lifecycle.
- **Amazon Elastic Compute Cloud (Amazon EC2)** - A web service that delivers secure, resizable compute capacity in the cloud.
- **Amazon Elastic Container Service for Kubernetes (Amazon EKS)** - Makes it easy to deploy, manage, and scale containerized applications using Kubernetes on AWS.
- **Amazon Elastic File System (Amazon EFS)** - Provides the company with a simple, scalable, fully managed elastic network file system (NFS) for use with AWS Cloud services and on-premises resources.
- **Amazon Elasticsearch Service** - A fully managed service that makes it easy for the company to deploy, secure, and operate Elasticsearch at scale with zero downtime.
- **Amazon Route 53** - A highly available and scalable cloud Domain Name System (DNS) web service that provides a reliable and cost-effective way to route the company's end users to internet applications.
- **Amazon Simple Storage Service (Amazon S3)** - A flexible way to store and retrieve data, providing the company with cost optimization, access control, and compliance.
- **Amazon Virtual Private Cloud (Amazon VPC)** - Enables the company to provision a logically isolated section of AWS where they can launch AWS resources in a virtual network that they define.
- **AWS Application Load Balancer (AWS ALB)** - To support content-based routing and applications that run in containers.
- **AWS Auto Scaling** - Monitors the company's applications and automatically adjusts capacity to maintain steady, predictable performance at the lowest possible cost.
- **AWS Certificate Manager** - Provisions, manages, and deploys public and private Secure Sockets Layer/Transport Layer Security (SSL/TLS) certificates to secure network communications and establish the identity of websites over the internet as well as resources on private networks.
- **AWS Identity and Access Management (AWS IAM)** - To control users' access to AWS services.
- **AWS Key Management Service (AWS KMS)** - A managed service that makes it easy for the company to create and control the encryption keys used to encrypt its data.
- **AWS Secrets Manager** - Enables the company to protect secrets needed to access its applications, services, and IT resources.
- **AWS Systems Manager Parameter Store** - Provides the company with secure, hierarchical storage for configuration data management and secrets management.
- **AWS Transit Gateway** - A service that enables the company to connect its Amazon VPCs and on-premises networks to a single gateway.



Global Energy & Petrochemical Leader's solution stack also includes additional, essential third-party tools:

- **Apache Airflow** - An open-source platform to programmatically author, schedule, and monitor workflows.
- **GitLab** - A complete DevOps platform for the entire software development and operations lifecycle, delivered as a single application.
- **HashiCorp Terraform** - An open-source tool that codifies APIs into declarative configuration files to enable the company to safely and predictably create, change, and improve infrastructure. nClouds used nCodeLibrary modules for infrastructure as code (IaC).
- **Helm** - Helps the company manage Kubernetes applications.
- **Istio** - An open-source Kubernetes service mesh that provides the fundamentals the company needs to successfully run a distributed microservice architecture.

- **Kubeflow** - A machine learning (ML) toolkit that makes deployments of ML workflows on Kubernetes simple, portable, and scalable.
- **Kubernetes Metrics Server** - A cluster-wide aggregator of resource usage data.
- **Ping Intelligent Identity Platform** - Provides intelligent access for customers, employees, and partners so they can securely connect to cloud, mobile, SaaS, and on-premises applications and APIs.

nClouds' Solution Architecture for Global Energy & Petrochemical Leader

nClouds helped Global Energy & Petrochemical Leader automate and deploy its next-generation energy platform regionally and globally. The project's scope was a high availability (HA), multi-tenant platform on AWS that enables users (the company's clients and trading partners) to access data lakes and process data related to its pipeline and supply chain.

The company collaborated with nClouds to implement AWS best practices to accelerate DevOps and enhance scalability. nClouds applied its expertise in DevOps, containerization, and migration.

nClouds delivered rich DevOps capabilities, including a CI/CD pipeline, infrastructure automation, and infrastructure as code (IaC) as part of building a well-architected new energy platform.

Infrastructure automation to accelerate DevOps includes:

- Terraform core infrastructure buildout and a CI/CD pipeline to automate deployment leveraging Amazon EKS, Jenkins, Apache Airflow, and Kubeflow.
- Amazon CloudWatch monitoring to automatically react to changes in AWS resources.
- AWS Auto Scaling to monitor applications and automatically adjust capacity.

For networking, nClouds implemented:

- A multi-tenant network.
- VPC peering connection between the engineering and data science Amazon VPCs.
- Istio Kubernetes service mesh.
- Network encryption of data at rest and data in transit.

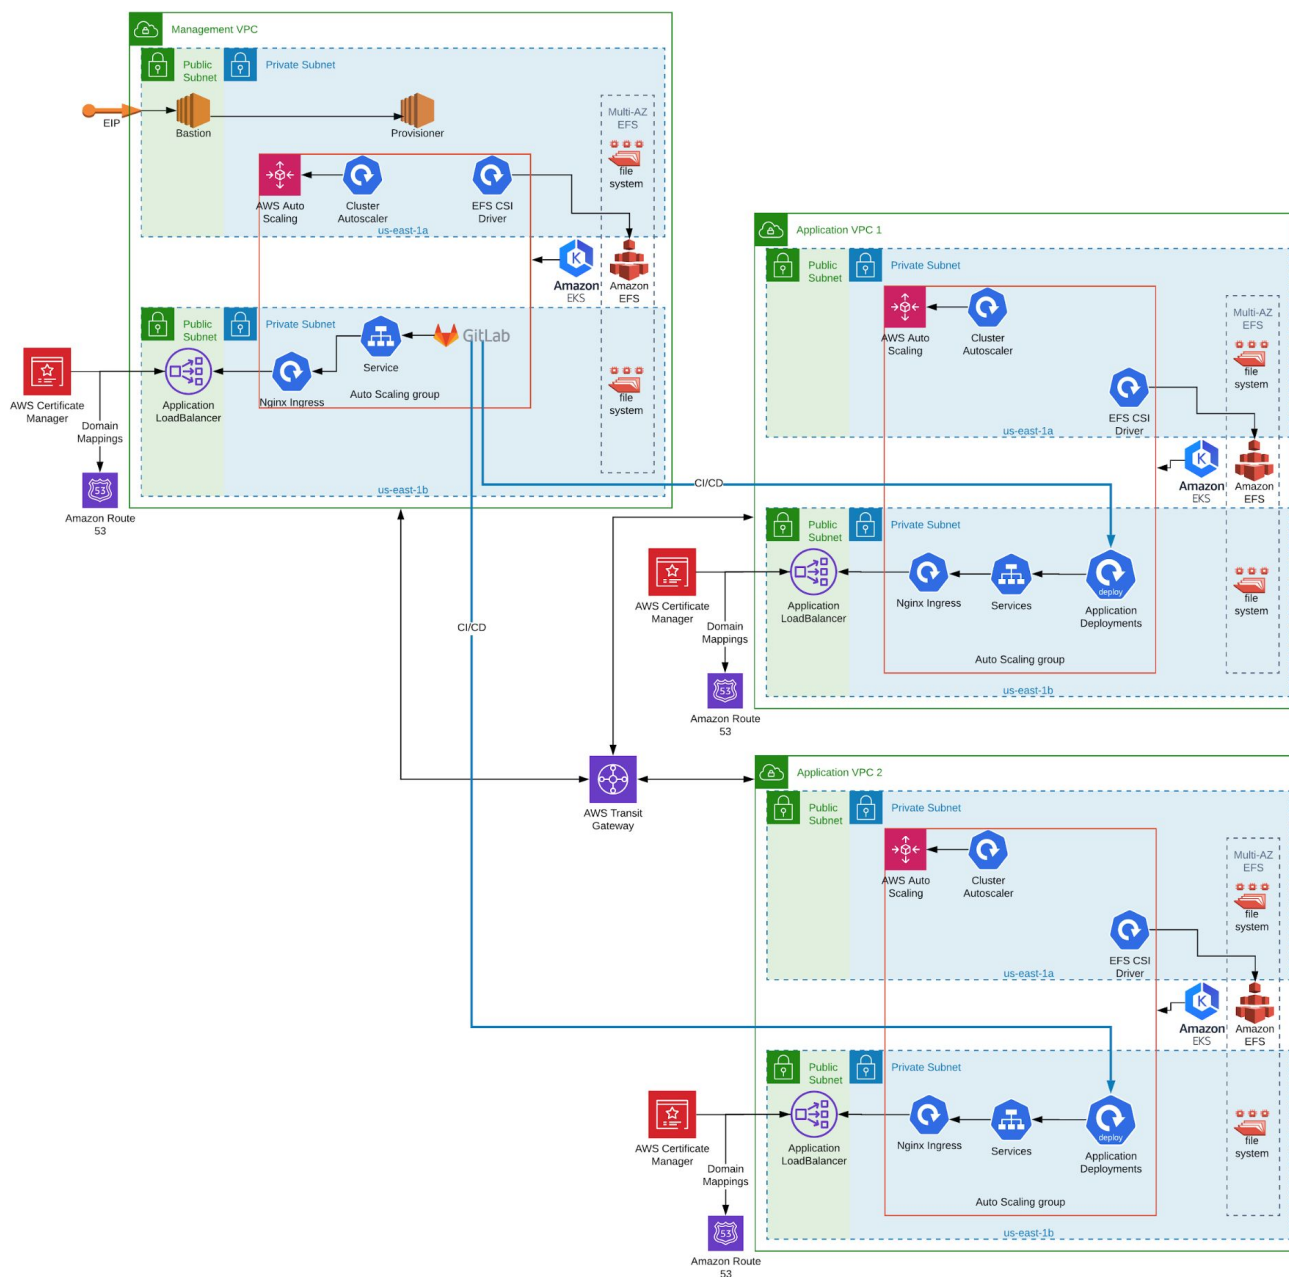
nClouds created three Amazon VPCs within the company's AWS account — a management VPC, a data engineering VPC, and a data science VPC.

- The **management VPC** contains the bastion host, a traffic and security gateway.
- The **data engineering and data science VPCs** have a public subnet and a database subnet in two Availability Zones to eliminate a single point of failure.
 - In one of the private subnets, Kubernetes Metrics Server connects to AWS Auto Scaling, and Amazon EFS Container Storage Interface (CSI) Driver connects to the company's file system.
 - In the other private subnet are Jenkins Deployment, Jenkins Service, and AWS ALB Ingress Controller, which connect via AWS ALB Ingress (in one of the public subnets) to the Jenkins Domain in the Amazon Cloud. The Jenkins Domain and an External DNS Controller (in the private subnet) connect to Amazon Route 53.
 - AWS Transit Gateway (which connects to the Amazon VPCs and the company's on-premises networks) and an AWS Certificate Manager (which connects to AWS ALB Ingress in the public subnet).
 - Amazon EFS and Amazon RDS are set up in the private subnets across multiple Availability Zones.
 - Amazon EKS runs the Kubernetes management infrastructure across multiple AWS Availability Zones to eliminate a single point of failure.
- In the **data engineering VPC**, nClouds implemented containerization, monitoring, and routing to:
 - Store, manage, and deploy Docker container images (Amazon ECR).

- Deploy, manage, and scale containerized applications using Kubernetes (Amazon EKS running Python, GraphQL to build and query APIs, an S3 bucket for data ingestion, and Apache AirFlow to programmatically author, schedule, and monitor workflows).
- Monitor performance (Amazon CloudWatch).
- Route end users to internet applications (Amazon Route 53).
- In the **data science VPC**, nClouds implemented AWS IAM for enhanced security.


High-level architecture diagram of the Management VPC and the Data Engineering VPC:


(A Data Science VPC was also built with a structure similar to the Data Engineering VPC shown below.)



The Benefits

Teaming with nClouds on this project, Global Energy & Petrochemical Leader achieved these benefits:

 **Accelerated DevOps**
nClouds built and deployed the new architecture using automation to accelerate DevOps so that Global Energy & Petrochemical Leader will benefit from a shorter development cycle and increased release velocity – enabling faster innovation.

 **Enhanced scalability**
To enhance the scalability of Global Energy & Petrochemical Leader's new machine learning (ML) platform for its customers, nClouds implemented Amazon EKS, Amazon EFS, Amazon Elasticsearch Service, Amazon Route 53, AWS Auto Scaling, and Kubeflow.

About nClouds

nClouds is a certified, award-winning provider of AWS and DevOps consulting and implementation services. We partner with our customers, as extensions of their teams, to build and manage modern infrastructure solutions that deliver innovation faster. We leap beyond the status quo.

Copyright © 2022 nClouds, Inc. All rights reserved

