



» CUSTOMER CASE STUDY

Visualizing Data with Containerized Applications on Amazon EKS

Deploying containerized data visualization applications
on Amazon EKS



Executive Summary

Harnessing their expertise in AWS and software engineering, the Nebulaworks team meticulously crafted a set of standards for utilizing Kubernetes within the AWS environment to enable the use of a data management solution for Gilead Sciences. By incorporating industry cloud best practices and leveraging infrastructure-as-code principles, the team developed robust deployment strategies for their containerized applications and Kubernetes clusters using Terraform. These standardized approaches not only addressed the lack of a standard deployment method but also provided a reliable framework for the company to seamlessly integrate data tools.

About the Customer

Gilead Sciences is a renowned biopharmaceutical company that specializes in developing innovative therapies to address critical medical needs. Established in 1987, Gilead has emerged as a global leader in the healthcare industry, committed to advancing medical science and improving patient well-being. With a wide range of products and solutions, Gilead is dedicated to transforming lives through its groundbreaking advancements. The company recognizes the immense potential of technology in driving advancements within the healthcare sector. Gilead harnesses advanced data analytics, artificial intelligence (AI), and machine learning (ML) techniques to extract valuable insights from extensive healthcare data. These insights contribute to their research and development efforts, facilitating the discovery of new therapeutic targets and the creation of more precise and personalized treatment options.

The Challenge

The company faced a significant challenge in utilizing enterprise-grade data visualization and data access management tools to analyze their extensive data lake within AWS. The absence of a standardized approach for deploying solutions to their cloud platform hindered their ability to leverage these powerful tools effectively. Recognizing the critical importance of accelerating the discovery of medicinal therapies, the company sought the expertise of a skilled Nebulaworks software engineering team to address this issue.

Why AWS

With a comprehensive suite of cloud services, AWS enables efficient scalability, enhanced security, and cost optimization, providing immense value to the business. The global infrastructure

ensures high availability and reliable performance. The diverse range of services, from compute to storage, analytics, and machine learning, enables rapid innovation and quick deployment. AWS's robust security measures and compliance certifications ensure data protection. Our customer, who has approved AWS as a vendor, benefits from these advantages. By leveraging AWS, Gilead can focus on core competencies, drive digital transformation, and stay competitive.

Why Nebulaworks?

Gilead Sciences required an embedded DevOps team that had a breadth of experience with various cloud native technologies, and operated like a modern software engineering team. Nebulaworks has expertise in building cloud native applications using Infrastructure as Code (IAC), extensive experience with container based workloads and orchestrators including Kubernetes since 2014, and Continuous Integration/Continuous Delivery/Continuous Deployment (CI/CD1/CD2) pipeline development, management and integration. Since the Data Visualization tool required a Kubernetes cluster, the Nebulaworks engineering team introduced mature practices on managing Kubernetes applications and clusters using IAC in the cloud.

The Solution

Gilead desired to leverage Starburst, a third-party enterprise-grade data visualization and access management solution for accessing data within their data lake on AWS. The vendor required that the application be deployed on a Kubernetes cluster and provided Nebulaworks with Helm charts for deployment and installation of the tool and required resources. The team decided to leverage EKS due to its built-in management capabilities (i.e. unhealthy node management, ease of Kubernetes upgrades, management of control plane resources) as opposed to deploying and managing our own cluster using raw EC2 instances. The application workloads were run on EC2 instances due to large compute and memory requirements.

The team leveraged Terraform, an IAC tool to define the required EKS, EC2, IAM and RDS resources. Extensive use of Terraform's module capabilities enabled the team to create reusable infrastructure components that could be leveraged across multiple projects. Below is a list of some of the modules created specifically for the deployment:

- EKS module: A high-level module for deployment of an EKS cluster that supports fargate and EC2 worker node groups.
- EKS nodegroup module : Module for node group definitions that supports the use of custom AMIs.
- RDS module: Module that supports creation/configuration of an AWS Aurora DB cluster.

- Starburst module: High-level application module that defined all resources that the Starburst application required. This module includes references to the above-mentioned modules plus definitions of standalone AWS resources.

In addition to deployment of the Starburst tool, the Nebulaworks team was tasked with developing a system that would ensure that Amazon Machine Images (AMIs) used for EKS worker nodes met Gilead's minimum security baseline (MSB). A requirement that Gilead's security division had in place in order to leverage EKS for the Organization as a whole. The MSB included OS-level configuration as well as installation of security tools. The Nebulaworks team developed a CI/CD pipeline that made use of Github Actions, Ansible, Packer and AWS Lambda to automatically build, configure, test and share the AMIs to various AWS accounts within the AWS organization.

Results and Benefits

Through their diligent efforts, the Nebulaworks software engineering team empowered the company to leverage the full potential of their data lake. By streamlining the deployment process and ensuring consistency across environments, the team's solution accelerated the discovery of medicinal therapies. Their technical acumen and dedication to excellence played a pivotal role in enabling the company to make significant strides towards their goal of advancing healthcare through innovative solutions.

1.) A standardized EKS deployment mechanism was established, allowing the Gilead Cloud Infrastructure team to support team-by-team requests. After the creation of the Infrastructure as Code (IAC) for EKS cluster and node group deployments, the customer gained increased development speed. The number of EKS clusters dedicated to distinct teams rose from 0 to 6. Each team was provided with permissions to interact with a development, test, and production cluster to enable a release engineering process for the infrastructure and Kubernetes application deployments.

2.) An additional increase of speed came from automating the creation of the EKS Node Group AMIs. The client required specialized security tooling installed on all EC2 instances in AWS and the traditional installation path was to engage a Security Engineer to perform the installation which took multiple days to a week for the request to be fulfilled. Instead of relying on the Security Team to manually install the security agents, Nebulaworks leveraged Hashicorp's Packer server templating tool to expedite and streamline the installation resulting in faster deployments using IAC best practices.

3.) Reproducibility was also gained by using an IAC approach to all permissions and resources that the applications required. The client gained the ability to create additional EKS clusters as needed.