

Terraform-Driven Kubernetes Cluster Management in AWS

Abhiram Reddy Peddireddy*

Citation: Peddireddy AR. Terraform-Driven Kubernetes Cluster Management in AWS. *J Artif Intell Mach Learn & Data Sci* 2024, 2(1), 742-746. DOI: doi.org/10.51219/JAIMLD/abhiram-reddy-peddireddy/185

Received: 03 January, 2024; **Accepted:** 28 January, 2024; **Published:** 30 January, 2024

*Corresponding author: Abhiram Reddy Peddireddy, Flexera, USA, E-mail: abhiramreddy2848@gmail.com

Copyright: © 2024 Peddireddy AR., This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

A B S T R A C T

The paper thoroughly explores how Terraform, an Infrastructure, as Code (IaC) tool can be effectively used to automate, scale and oversee Kubernetes clusters within the AWS environment. It carefully analyzes the benefits of utilizing Terraform emphasizing its effects, on improving efficiency, automation, scalability and security in overseeing Kubernetes clusters. The study compares Terraform with approaches and other common IaC tools to showcase Terraforms capabilities in managing infrastructure. Additionally it delves into how Terraform integrates with AWS services to simplify processes and reduce complexities. The paper also discusses trends and potential advancements in integrating Kubernetes and Terraform to enhance the management of cloud native applications.

Keywords: Infrastructure as Code (IaC), Kubernetes, AWS, Terraform, Automation, Scalability, Security, Cloud Infrastructure, DevOps, Cloud-native Applications

1. Introduction

Kubernetes, also known as K8s is an open source platform used for automating the deployment, scaling and operation of application containers. It plays a role, in application deployment by effectively managing containerized applications across multiple machines within a cluster. This results in enhanced scalability, reliability and efficiency. Kubernetes orchestrates the running of containerized applications on a cluster by handling tasks like load balancing, resource allocation and ensuring operation of application instances¹.

Amazon Elastic Kubernetes Service (EKS) is a managed service provided by Amazon that simplifies the process of running Kubernetes on AWS. With EKS users do not have to worry about setting up and maintaining their Kubernetes control plane as it is automatically managed. EKS takes care of tasks like upgrades and patching to ensure availability and security. By allowing AWS to manage the Kubernetes control plane developers can concentrate on developing their applications and making use of integrations with AWS services¹¹.

Terraform is an open source tool for Infrastructure as Code (IaC) created by HashiCorp. It enables users to define and provision data center infrastructure using HashiCorp Configuration Language (HCL) or JSON- two high level configuration languages. Terraform facilitates automation in infrastructure management, by promoting reproducibility and minimizing errors through its configuration approach. Using terraform to manage kubernetes clusters on AWS offers advantage. It allows organizations to define their infrastructure and application deployment processes, in code promoting collaboration, consistency and automation in line with DevOps principles. Terraform simplifies the setup and maintenance of EKS clusters by providing modules and plugins that seamlessly integrate with AWS services like IAM, VPC and auto scaling groups¹⁰.

The synergy, between Kubernetes orchestration capabilities and Terraforms infrastructure automation boosts the efficiency, scalability and maintainability of applications.

1.1. Objective

The main goal of the paper is to demonstrate how Terraform, an Infrastructure, as Code (IaC) tool can effectively automate, scale and oversee Kubernetes clusters within the AWS environment. The paper carefully examines the benefits of using Terraform highlighting its impact on improving efficiency, automation, scalability and security in managing Kubernetes clusters.

Through a comparison with methods and other common IaC tools the paper aims to provide a thorough analysis showcasing Terraform's superior capabilities in managing infrastructure. By focusing on the nature of Terraform and its integration with AWS services the paper intends to illustrate the processes and reduced operational complexities achieved through this approach.

Furthermore the paper delves into trends and potential advancements in integrating Kubernetes and Terraform. It explores developments that could enhance the smoothness and effectiveness of managing native applications contributing to the continuous evolution of infrastructure as code and cloud management practices.

In summary the paper seeks to offer an in depth assessment of how Terraform's reshaping Kubernetes cluster management on AWS offering insights, for professionals and researchers involved in cloud infrastructure management⁷.

2. Literature Review

2.1. Terraform and Infrastructure as Code

Terraform, created by HashiCorp has become a tool, in the field of Infrastructure as Code (IaC). It allows for defining, provisioning and managing cloud infrastructure using a language that's easy to understand. The key role of Terraform in managing infrastructure lies in its capability to establish a repeatable process for setting up and handling infrastructure across cloud platforms and service providers. By translating infrastructure into code Terraform enables the automation of environment setup reducing the chances of errors enhancing efficiency and enabling version control over infrastructure configurations¹¹.

An essential aspect of Terraform is its state management system, which monitors the status of the infrastructure⁸. This feature helps Terraform identify changes to achieve the desired configuration outlined in the files. This method simplifies handling environments and supports adopting practices in DevOps, like continuous integration and continuous deployment (CI/CD).

2.1.1. Comparison with Other IaC Tools: When we look at Infrastructure, as Code (IaC) tools like AWS Cloud Formation and Ansible Terraform stands out with its advantages. Unlike Cloud Formation, which is limited to AWS Terraform supports a variety of cloud providers such as AWS, Google Cloud Platform and Microsoft Azure making it a versatile option for cloud strategies⁴. Moreover Terraform's modularity enables the creation of configurations simplifying the management of infrastructure setups⁵. While Ansible focuses more on configuration management. Can handle tasks beyond provisioning infrastructure Terraform excels, in defining and managing immutable infrastructure. This makes it especially well-suited for creating and overseeing cloud resources from the

ground up⁸.

2.2. Kubernetes cluster management

2.2.1. Review of traditional methods for managing Kubernetes clusters: In the sense handling Kubernetes clusters typically involves utilizing command line tools like kubectl, for orchestrating clusters alongside scripts and configuration files to oversee deployments, services and networking. While these methods offer a level of control and adaptability they can be intricate and prone to errors in expansive environments. Manual cluster management often demands expertise. Can result in inconsistencies and challenges in maintaining infrastructure as it expands¹³.

Additionally conventional approaches may lack integration with CI/CD pipelines making it more challenging to automate deployments and efficiently manage application lifecycles. Although tools such as Helm have emerged to simplify package management in Kubernetes they do not fully tackle the complexities of overseeing the underlying infrastructure.

2.3. Advantages of using terraform for Kubernetes cluster management on AWS

Harnessing Terraform for managing Kubernetes clusters¹³ on AWS presents advantages. To begin with Terraform's declarative methodology enables users to define the infrastructure. Including EKS clusters in code. This facilitates version control, auditing and collaboration processes while ensuring that changes, to infrastructure are monitored and reproducible. Additionally Terraform's wide range of modules, like the AWS EKS module makes setting up Kubernetes clusters easier by offering to use templates that follow industry practices. Another key benefit is how seamlessly Terraform can be integrated with AWS services. For example it can handle IAM roles, VPC configurations and auto scaling groups to ensure that the Kubernetes cluster is not just deployed correctly but smoothly connected to the AWS environment for improved security and scalability. This comprehensive method of managing infrastructure reduces complexities. Enhances the dependability and performance of Kubernetes clusters^{2,3}. To sum up Terraform offers an automated solution for overseeing Kubernetes clusters, on AWS overcoming the constraints of traditional approaches and providing a more efficient way to manage infrastructure through code⁶.

3. Methodology

3.1. Setting up terraform for AWS

- 1. Installation and configuration of terraform:** The initial step in leveraging Terraform for Kubernetes cluster management on AWS involves installing Terraform. This process starts with downloading the Terraform binary from the official HashiCorp website, verifying its integrity, and adding it to the system's PATH. Once installed, a working directory for Terraform scripts is created and initialized using the terraform init command, which prepares the directory by downloading necessary provider plugins and setting up the backend for storing state files.
- 2. Setting up AWS credentials and configuring the AWS provider in terraform:** To enable Terraform to interact with AWS, it is necessary to configure AWS credentials. This involves creating an IAM user with the required permissions and configuring the AWS CLI with these

credentials. The credentials are stored in a configuration file that Terraform uses for authentication. In the Terraform configuration file, the AWS provider is specified with the appropriate region and credentials profile, ensuring secure and efficient management of AWS resources.

3.2. Provisioning Kubernetes Clusters

1. Using the AWS EKS Module to Define the Desired State of the Cluster: The AWS EKS module simplifies the process of provisioning Kubernetes clusters. This module encapsulates best practices and provides a reusable template for creating EKS clusters. The desired state of the cluster, including the Kubernetes version, VPC CIDR block, and availability zones, is defined in the Terraform configuration file. This setup ensures that the cluster is created according to specified requirements.

2. Defining Subnets, Node Groups, and IAM Roles: Defining subnets, node groups, and IAM roles is crucial for setting up the Kubernetes cluster. Subnets provide the necessary networking infrastructure, node groups define the instances that will run Kubernetes workloads, and IAM roles ensure secure access to AWS resources. The configuration includes specifying the number and type of instances for node groups, as well as mapping IAM roles for user access and permissions.

3. Execution of Terraform Commands to Provision Resources: Once the configuration is defined, the provisioning process is initiated with Terraform commands. The terraform init command initializes the configuration directory, downloading the necessary provider plugins. The terraform apply command is then used to apply the configuration and create the resources. This command displays a plan of the changes to be made and prompts for confirmation before proceeding. Upon confirmation, Terraform provisions the EKS cluster and associated resources, providing real-time updates and the final state upon completion.

3.3. Managing Kubernetes resources with terraform

1. Configuring the Kubernetes provider in terraform: After the EKS cluster is operational, Terraform can be used to manage Kubernetes resources. This involves configuring the Kubernetes provider in Terraform to interact with the Kubernetes API. The provider configuration includes details such as the cluster endpoint, CA certificate, and authentication token, enabling Terraform to manage resources within the cluster.

2. Managing Kubernetes resources like pods, services, and deployments: Terraform scripts are used to define and manage Kubernetes resources such as pods, services, and deployments. For instance, a deployment for an application can be defined with specific metadata, replica count, selectors, and container specifications. Terraform manages the lifecycle of these resources, ensuring they are deployed and maintained according to the specified configuration.

In conclusion, using Terraform for managing Kubernetes clusters on AWS involves detailed steps for setting up Terraform, configuring AWS, provisioning an EKS cluster, and managing Kubernetes resources with Terraform scripts. This methodology enhances automation, consistency, and scalability in managing cloud-native applications.

```

1 # Setting Up Terraform for AWS
2 provider "aws" {
3   region = "us-west-2"
4 }
5
6 # Provisioning Kubernetes Clusters
7 module "eksctl-cluster" {
8   source = "git::github.com:terraform-aws-modules/eks-cluster"
9 }
10
11 providers {
12   aws = aws
13   aws.shared = aws.shared
14 }
15
16 AWS_REGION = "us-west-2"
17 ENVIRONMENT = "dev"
18 EKS_CLUSTER_KUBERNETES_VERSION = "1.25"
19
20 vpc_cidr = "19.28.0.0/16"
21 eksctl_gateway_id = "tgw-1234567890abcdef"
22 vpc_availability_zone_ids = ["us-west-2a", "us-west-2b", "us-west-2c"]
23
24 # Node Group Configuration
25 CLUSTER_NODE_SELECTED_SIZE = 3
26 CLUSTER_NODE_MIN_SIZE = 3
27 CLUSTER_NODE_MAX_SIZE = 3
28 CLUSTER_NODE_INSTANCE_TYPE = "m5.large"
29
30 # Additional Settings
31 EKSC2_CLUSTER_NODE_SELECTED_SIZE = 0
32 EKSC2_CLUSTER_NODE_MIN_SIZE = 0
33 EKSC2_CLUSTER_NODE_MAX_SIZE = 0
34 DISABLE_OCI = var.DISABLE_OCI
35
36 INSTALL_AWS_IAM_INGRESS_IAM_POLICIES = true
37 INSTALL_IK_CONTROLLER = true
38
39 # Custom security group rules for worker nodes
40 worker_custom_security_group_rules = [
41   kong_https443 = {
42     type = "ingress"
43     from_port = 443
44     to_port = 443
45     protocol = "tcp"
46     cidr_blocks = ["0.0.0.0/0"]
47   },
48   kong_https1443 = {
49     type = "ingress"
50     from_port = 1443
51     to_port = 1443
52     protocol = "tcp"
53     cidr_blocks = ["0.0.0.0/0"]
54   },
55   kong_https2443 = {
56     type = "ingress"
57     from_port = 2443
58     to_port = 2443
59     protocol = "tcp"
60     cidr_blocks = ["0.0.0.0/0"]
61   },
62 ]
63
64 additional_vpc_cidrs = ["18.248.16.0/22", "18.248.32.0/22"]
65 additional_private_subnet_cidrs = ["18.248.16.0/22", "18.248.26.0/22", "18.248.32.0/22", "18.248.36.0/22", "18.248.48.0/22"]
66
67 # IAM Roles
68 map_roles = [
69   {
70     role_name = "arn:aws:iam::123456789012:role/Admin"
71     groups = ["system-masters"]
72   },
73   {
74     role_name = "arn:aws:iam::123456789012:role/aws-reserved/sso.amazonaws.com/AWSReservedSSO_AdministratorAccess_1234567890"
75     groups = ["iam-cluster-reader-group"]
76   },
77 ]
78
79 # Managing Kubernetes Resources with Terraform
80 provider "kubernetes" {
81   host = module.eksctl-cluster.cluster_endpoint
82   cluster_ca_certificate = base64decode(module.eksctl-cluster.cluster_certificate_authority_data)
83   token = data.aws_iam_role_arn.cluster_arn.token
84 }
85
86 resource "kubernetes_deployment" "nginx" {
87   metadata {
88     name = "nginx-deployment"
89   }
90   spec {
91     replicas = 2
92     selector {
93       match_labels = {
94         app = "nginx"
95       }
96     }
97     template {
98       metadata {
99         labels = {
100           app = "nginx"
101         }
102       }
103       spec {
104         container {
105           name = "nginx"
106           image = "nginx:1.14.2"
107           ports {
108             container_port = 80
109           }
110         }
111       }
112     }
113   }
114 }
115
116 }

```

Figure 1: An example terraform workflow for provisioning and managing Kubernetes clusters on AWS.

4. Results and Discussion

4.1. Provisioning Efficiency

1. Evaluation of the efficiency and speed of provisioning Kubernetes clusters using terraform: The evaluation of provisioning efficiency involves measuring the time and resources required to set up Kubernetes clusters using Terraform. Terraform's declarative approach allows for defining the desired state of infrastructure, which Terraform then reconciles with the actual state. This process is highly efficient, as Terraform automates the creation, modification, and deletion of infrastructure components based on the configuration files. The speed of provisioning is significantly enhanced due to Terraform's ability to parallelize resource creation where dependencies allow, reducing the overall time required for deployment.

2. Comparison with Traditional Methods of Cluster Management: Traditional methods of managing Kubernetes clusters often involve manual configuration and scripting. This approach can be time-consuming and prone to human error. In contrast, Terraform provides a more streamlined and automated process. By using Infrastructure

as Code (IaC), Terraform ensures that infrastructure setups are consistent and repeatable. The traditional methods lack the scalability and repeatability offered by Terraform, making Terraform a superior choice for efficient and rapid provisioning of Kubernetes clusters⁷.

4.2. Automation and Scalability

1. **Benefits of automation in cluster management using terraform:** Automation in cluster management using Terraform offers numerous benefits, including reduced manual intervention, minimized errors, and consistent environments. Terraform scripts can be reused and version-controlled, enabling teams to manage infrastructure changes collaboratively and efficiently. The automation capabilities of Terraform extend to scaling operations, allowing for the dynamic adjustment of resources based on demand. This automated scaling ensures that the cluster⁷ can handle varying workloads without manual reconfiguration.
2. **Analysis of scalability and performance improvements in managing Kubernetes clusters:** Scalability is a critical aspect of managing Kubernetes clusters, and Terraform excels in this area. Terraform's modularity and use of IaC principles facilitate the scalable management of large, complex environments. Performance improvements are realized through the efficient allocation and management of resources, as Terraform optimizes the infrastructure based on the defined configurations¹². The ability to automate scaling operations ensures that resources are used optimally, enhancing both performance and cost efficiency.

4.3. Security and Compliance

1. **Assessment of Security Best Practices Implemented via Terraform:** Terraform allows for the implementation of security best practices through its configuration files. These practices include defining secure access controls, encrypting data in transit and at rest, and ensuring compliance with organizational policies. By codifying security configurations, Terraform ensures that all infrastructure components adhere to the specified security standards, reducing the risk of misconfigurations and vulnerabilities.
2. **Evaluation of Compliance Management for Kubernetes Clusters on AWS:** Compliance management is critical for organizations operating in regulated industries. Terraform aids in compliance by providing a clear and auditable trail of infrastructure changes. The configuration files serve as documentation of the infrastructure setup, which can be reviewed and audited to ensure compliance with industry standards and regulations. Additionally, Terraform's integration with AWS services enables the use of AWS compliance tools and features, further enhancing the ability to manage compliance effectively.

In summary, the use of Terraform for provisioning and managing Kubernetes clusters on AWS demonstrates significant improvements in efficiency, automation, scalability, security, and compliance. These enhancements make Terraform an invaluable tool for modern infrastructure management, offering a robust solution for deploying and maintaining Kubernetes clusters in a consistent, repeatable, and secure manner.

5. Conclusion and Future Scope

5.1. Summary of findings

1. **Advantages of using terraform for Kubernetes cluster**

management in AWS: Managing Kubernetes clusters on AWS using Terraform comes with a range of benefits. Terraform's method of Infrastructure as Code (IaC) makes it easier to handle infrastructures by letting users specify the desired state of resources, which Terraform then aligns with the state. This method ensures uniformity, consistency and promotes team-work collaboration. Moreover Terraform smoothly integrates with AWS services utilizing the cloud platforms features to improve infrastructure management.

2. Summary of improvements in efficiency, automation, scalability, and security:

1. **Efficiency:** Terraform streamlines the setup of infrastructure saving a lot of time and effort compared to doing it. This automation makes sure that resources are created, changed and removed efficiently based on the specified configurations.
2. **Automation:** Terraform's automation features go beyond setting up infrastructure to include tasks, like adjusting scale and updating resources over time. This decreases the need for work reduces errors and ensures that the infrastructure stays as intended.
3. **Scalability:** With its design and Infrastructure as Code (IaC) principles Terraform makes it easy to manage infrastructure at scale. Organizations can quickly adapt resources to meet varying needs guaranteeing performance and cost effectiveness.
4. **Security:** By allowing security best practices to be written into code Terraform ensures that all parts of the infrastructure comply, with rules and industry standards. This lowers the chances of misconfigurations and boosts security measures.

5.2. Future trends

1. **Potential developments in Kubernetes and terraform integration:** As both Kubernetes and Terraform continue to evolve, their integration is expected to become even more seamless and powerful. Future developments may include enhanced support for multi-cloud environments, allowing organizations to manage Kubernetes clusters across different cloud providers with a unified approach. Additionally, advancements in Terraform's ecosystem, such as new modules and providers, will likely provide even more robust solutions for Kubernetes management.
2. **Future research directions and scope for improvement in IaC tools and Kubernetes management:**
 1. **Enhanced automation and AI integration:** Future research may explore the integration of artificial intelligence and machine learning with IaC tools like Terraform. This could lead to smarter automation, predictive scaling, and more efficient resource management.
 2. **Security and compliance automation:** As security and compliance requirements become more stringent, there is a growing need for tools that can automatically enforce policies and ensure compliance. Research in this area could lead to the development of advanced features that simplify compliance management.
 3. **User experience and collaboration:** Improving the user experience of IaC tools and enhancing collaboration capabilities will be crucial. Future enhancements could include more intuitive interfaces, better collaboration

features, and improved version control systems.

4. **Performance optimization:** Ongoing research into the performance optimization of Kubernetes clusters and the underlying infrastructure will be important. This could involve developing new techniques for resource allocation, load balancing, and minimizing latency⁹.

6. Conclusion

In conclusion, Terraform has proven to be a valuable tool for managing Kubernetes clusters on AWS, offering significant improvements in efficiency, automation, scalability, and security. As technology continues to advance, the integration between Terraform and Kubernetes will likely become even more robust, driving further innovation in infrastructure management. Future research and development in this field hold the potential to revolutionize how organizations deploy and manage their cloud-native applications, ensuring they can meet the demands of an ever-changing technological landscape.

7. References

1. Gupta M, Sanjana K, Akhilesh K, Chowdary M. Deployment of multi-tier application on cloud and continuous monitoring using Kubernetes. 2021 5th International Conference on Electrical, Electronics, Communication, Computer Technologies and Optimization Techniques (ICEECCOT) 2021; 602-607.
2. Reddy Y, Reddy P, Ganesan N, Thangaraju B, Performance study of kubernetes cluster deployed on openstack, VMs and baremetal. 2022 IEEE International Conference on Electronics, Computing and Communication Technologies (CONECCT) 2022; 1-5.
3. Orzechowski M, Balis B, Pawlik K, Pawlik M, Malawski M. Transparent deployment of scientific workflows across clouds-Kubernetes approach. 2018 IEEE/ACM International conference on utility and cloud computing Companion (UCC Companion) 2018; 9-10.
4. Bahaweres RB, Najib FM. Provisioning of disaster recovery with terraform and Kubernetes: A case study on software defect prediction. 2023 10th International Conference on Electrical Engineering, Computer Science and Informatics (EECSI) 2023; 183-189.
5. Ekanayaka E, Thathsarani J, Karunanayaka D, Kuruwitaarachchi N, Skandakumar N. Enhancing devops infrastructure for efficient management of microservice applications. 2023 IEEE International Conference on e-Business Engineering (ICEBE) 2023; 63-68.
6. Sheikh S, Suganya G, Premalatha M. Automated resource management on AWS cloud platform. Proceedings of 6th International Conference on Big Data and Cloud Computing Challenges 2019.
7. Sindhu G, N. Mtech, and R. Pavithra D. Deploying a Kubernetes Cluster with Kubernetes Operation (kops) on AWS Cloud: Experiments and Lessons Learned. Int J Engineering Advanced Technology 2020.
8. Campbell B. Terraform In-Depth. The definitive guide of AWS infrastructure automation 2019; 123-203.
9. Sahana B, Kumaraswamy T, Nachiketh RG, Navadeep S, Noronha J. Weight based load balancing in Kubernetes using AWS. 2023 International Conference on Intelligent Data Communication Technologies and Internet of Things (IDCIoT) 2023; 629-634.
10. Bailuguttu S, Chavan A, Pal O, Sannakavalappa K, Chakrabarti D. Comparing performance of bastion host on cloud using Amazon web services vs terraform. Indonesian J Electrical Engineering Computer Science 2023.
11. Ganeshan M, Malathi S. Building and deploying a static application using Jenkins and Docker in AWS, Int J Trend in Scientific Research and Development 2020.
12. Haragi LD, Mahith S, Sahana B. Infrastructure Optimization in Kubernetes Cluster. J University of Shanghai for Science and Technology 2021.
13. Hui A, Lee B. Epsilon: A microservices based distributed scheduler for Kubernetes cluster. 2021 18th Int Joint Conference on Computer Science and Software Engineering (JCSSE) 2021; 1-6.