

Infrastructure as Code (IaC): Best Practices of Implementing IaC, Especially in Automating Infrastructure Provisioning and Management Using Terraform

Gowtham Mulpuri

Silicon Labs, TX, USA

Email: [gowtham.mulpuri\[at\]silabs.com](mailto:gowtham.mulpuri[at]silabs.com)

Abstract: *Infrastructure as Code (IaC) is a paradigm shift in IT infrastructure management, advocating for the automation and management of physical and virtual infrastructure through code. Terraform, a tool created by HashiCorp, has emerged as a leader in this space. It allows for the definition, provisioning, and management of infrastructure across a variety of cloud providers and services using a declarative configuration language. This paper aims to explore best practices in implementing IaC with an emphasis on Terraform, covering concepts from code organization to operational considerations. Infrastructure as Code'' (IaC) is a key practice in the field of DevOps and cloud computing, where the infrastructure is provisioned and managed using code instead of through manual processes. This approach enables developers and system administrators to automate the setup, deployment, and management of infrastructure in a consistent and repeatable way. Below, I'll detail some best practices for implementing IaC, particularly with a focus on Terraform. This paper discusses the concept of Infrastructure as Code (IaC), emphasizing its implementation using Terraform. It explores best practices for automating infrastructure provisioning and management, highlighting the advantages and challenges of IaC in modern cloud environments. The paper aims to guide IT professionals in effectively utilizing Terraform for scalable and efficient infrastructure deployment.*

Keywords: IaC, Terraform, HashiCorp, CICD Pipelines, Automation, Infrastructure Automation

1. Introduction

The traditional approach to infrastructure management involved manual processes which were error - prone, non - repeatable, and difficult to track. IaC introduces a systematic approach where infrastructure is provisioned and managed using code. This code is versioned, tested, and reused, leading to more consistent and reliable environments. Terraform's declarative syntax and provider ecosystem make it an ideal choice for implementing IaC.

1.1 Definition of IaC

Infrastructure as Code (IaC) is a method of managing and provisioning computing infrastructure through machine - readable definition files, rather than physical hardware configuration or interactive configuration tools. This approach enables developers and IT operations teams to automatically manage, monitor, and provision resources, rather than manually setting up and configuring hardware or virtual machines.

1.2 Importance of IaC:

1.2.1 Consistency and Standardization: IaC ensures that the same configuration is applied uniformly, reducing the likelihood of discrepancies or errors.

1.2.2 Speed and Efficiency: Automates the provisioning process, enabling rapid deployment of infrastructure, which is crucial in agile environments and for scaling operations.

1.2.3 Version Control and Documentation: Changes to the infrastructure are version - controlled, providing a clear audit trail and documentation.

1.2.4 Cost Reduction: Reduces the need for manual intervention, thereby cutting down on labor costs and minimizing the risk of human error.

1.2.5 Scalability and Flexibility: Facilitates easy scaling of infrastructure to meet changing demands and supports the management of a diverse and distributed infrastructure.

1.2.6 Disaster Recovery: Enhances the ability to quickly recover from failures and reduce downtime by enabling consistent redeployment of a specific infrastructure state.

1.3 Overview of Terraform

Terraform is an open - source infrastructure as code tool created by HashiCorp. It allows users to define and provision a datacenter infrastructure using a high - level configuration language known as HashiCorp Configuration Language (HCL), or optionally JSON.

1.3.1 Declarative Configuration: Terraform uses a declarative approach where users specify the desired end - state and Terraform works to achieve that state.

1.3.2 Provider Ecosystem: Supports multiple service providers (like AWS, Azure, Google Cloud, etc.), allowing users to manage a wide range of resources.

1.3.3 State Management: Terraform maintains a state file to keep track of the resources it manages, which aids in change management and resource mapping.

1.3.4 Modularity and Reusability: Supports modules for creating reusable components, promoting code reuse and maintainability.

1.3.5 Idempotency: Running the same configuration multiple times produces the same result, ensuring consistency and stability.

1.3.6 Plan and Apply Workflow: Terraform generates an execution plan before making any changes, allowing users to review what will be done before it is executed, enhancing control and visibility.

In summary, IaC, exemplified by tools like Terraform, represents a significant shift in the way infrastructure is managed, making it more consistent, efficient, and scalable. Terraform, with its declarative approach and extensive provider support, is a popular choice in this field.

2. Implementing IaC with Terraform:

2.1 Setting up Terraform:

Terraform, developed by HashiCorp, is an open - source infrastructure as code software tool that enables users to define and provision a datacenter infrastructure using a high - level configuration language known as HashiCorp Configuration Language (HCL), or optionally JSON. To begin using Terraform for infrastructure management, a basic setup involves the following steps:

2.1.1 Installation: Terraform is distributed as a single binary. Users need to download and install the Terraform binary from the official *website*¹

It is available for various platforms including Windows, macOS, and Linux.

2.1.2 Configuration: Terraform uses configuration files to describe the components needed to run a single application or your entire datacenter. These files end with .tf or .tf.json to indicate their HCL or JSON syntax, respectively.

2.1.3 Initialization: A Terraform configuration or project is initialized using the terraform init command. This command performs several different initialization steps to prepare a working directory for use.

2.2 Basic Terraform Operations:

2.2.1 Write: The first step in using Terraform is to write the infrastructure code. This code can be written in HCL or JSON.

2.2.2 Plan: Before applying the changes, Terraform performs a dry run to show what actions will be taken based on the current configuration.

2.2.3 Apply: If the plan is acceptable, the terraform apply command is used to execute the actions proposed in a Terraform plan.

```
provider "aws" {
  region = "us-west-2"
}

resource "aws_instance" "example" {
  ami          = "ami-0c55b159cbf1f0"
  instance_type = "t2.micro"
}
```

Figure 1: Terraform Provider

In this basic example listed in the Figure 1, Terraform is configured to provision an AWS EC2 instance in the US West (Oregon) region.

2.3 Advanced Terraform Features: Terraform also offers advanced features such as:

2.3.1 Workspaces: Used to manage multiple distinct sets of infrastructure resources or environments.

2.3.2 Remote Backends: For state management in a remote shared store, which is essential for team collaboration.

2.3.3 Dependency Management: Terraform automatically understands dependency relationships between resources.

3. Best Practices for IaC with Terraform:

3.1 Version Control

Keep all your Terraform configurations in a version control system (like Git). This practice not only tracks changes over time but also allows team collaboration.

3.2 Modularize Your Code

Structure your Terraform code into modules. Each module should represent a logical component of your infrastructure (like network, compute, storage) and can be reused in different environments or projects.

3.3 Use Remote State

Terraform stores state about your managed infrastructure and configuration. This state should be stored in a remote data store like Terraform Cloud or AWS S3 for team access and to avoid conflicts.

3.4 Immutable Infrastructure

Aim for an immutable infrastructure model where you replace instances rather than updating them. This approach reduces inconsistencies and potential errors during updates.

3.5 Environment Separation

Use different configurations or workspaces for different environments (development, staging, production) to avoid accidental changes in the production environment.

3.6 Automate Workflow

Integrate Terraform with your CI/CD pipeline for automated testing and deployment of your infrastructure changes.

3.7 Least Privilege Access

Implement least privilege access in your Terraform scripts. Ensure that the execution context (like a CI/CD pipeline) has only the necessary permissions to perform its tasks.

3.8 Code Review and Testing

Just like application code, Terraform code should go through code reviews and testing. Use tools like Terraform fmt and Terraform validate for formatting and validating your code, respectively.

3.9 Documentation

Document your Terraform configurations and modules. This documentation should explain the purpose, usage, and any variables or outputs of the modules.

3.10 Regularly Update and Refactor

Infrastructure needs change over time, and so should your Terraform code. Regularly update and refactor your code to align with best practices and new features of Terraform.

3.11 Disaster Recovery Plan

Always have a disaster recovery plan. Test your Terraform configurations for disaster recovery scenarios to ensure your infrastructure can be quickly recreated.

3.12 Monitoring and Logging

Implement monitoring and logging to track the performance and health of your infrastructure. This practice helps in proactive issue resolution and audit trails.

3.13 Advanced Terraform Practices:

- State Locking and State Environments:** Use state locking to prevent concurrent runs from causing conflicts. Also, consider using different state files for different environments to reduce blast radius in case of issues.
- Dynamic Configuration:** Use Terraform's dynamic blocks and other advanced features for more flexible and scalable infrastructure configurations.
- Integrate with Cloud - Native Tools:** For cloud environments, integrate Terraform with cloud - native tools and services for better resource management.

By following these best practices, you can effectively use Terraform to implement Infrastructure as Code, leading to more efficient, reliable, and scalable infrastructure management. Remember, the key to successful IaC is not just about the tools like Terraform, but also about the practices, processes, and collaboration within the team.

4. Modularization and Reusability

One of the key best practices in Terraform is to organize your infrastructure code into modules. Modules allow you to package and reuse code for similar infrastructure setups. For example, if you frequently deploy AWS EC2 instances with a similar configuration, you can create a module for that.

```
# File: modules/ec2/main.tf

variable "instance_type" {
  description = "EC2 instance type"
  type        = string
  default     = "t2.micro"
}

variable "ami_id" {
  description = "AMI ID for EC2 instance"
  type        = string
}

resource "aws_instance" "example" {
  ami           = var.ami_id
  instance_type = var.instance_type
  tags = {
    Name = "ExampleInstance"
  }
}
```

Figure 2: Module for AWS EC2 Instance

The above figure shows an example code for implementing AWS EC2 Instance using module concept.

5. Automation and CI/CD Integration:

Automating Terraform through CI/CD pipelines ensures consistent and error - free deployment of infrastructure.

Sample CI/CD Pipeline Steps:

Linting and Validation: Run terraform fmt and terraform validate.

Plan: Execute terraform plan in CI/CD and present the output for review.

Apply: After approval, run terraform apply to deploy the changes.

A typical CI/CD pipeline with Terraform integration might include the following stages:

Source: Code is written and committed to a version control system (e. g., Git).

Build: The application is compiled or packaged.

Test: Automated tests are run to ensure the application behaves as expected.

Plan (Terraform): Terraform's plan operation is executed to preview infrastructure changes.

Apply (Terraform): Terraform's apply operation is executed to update the infrastructure.

Deploy: The application is deployed to the updated infrastructure.

Monitor: The application and infrastructure are monitored for performance and reliability.

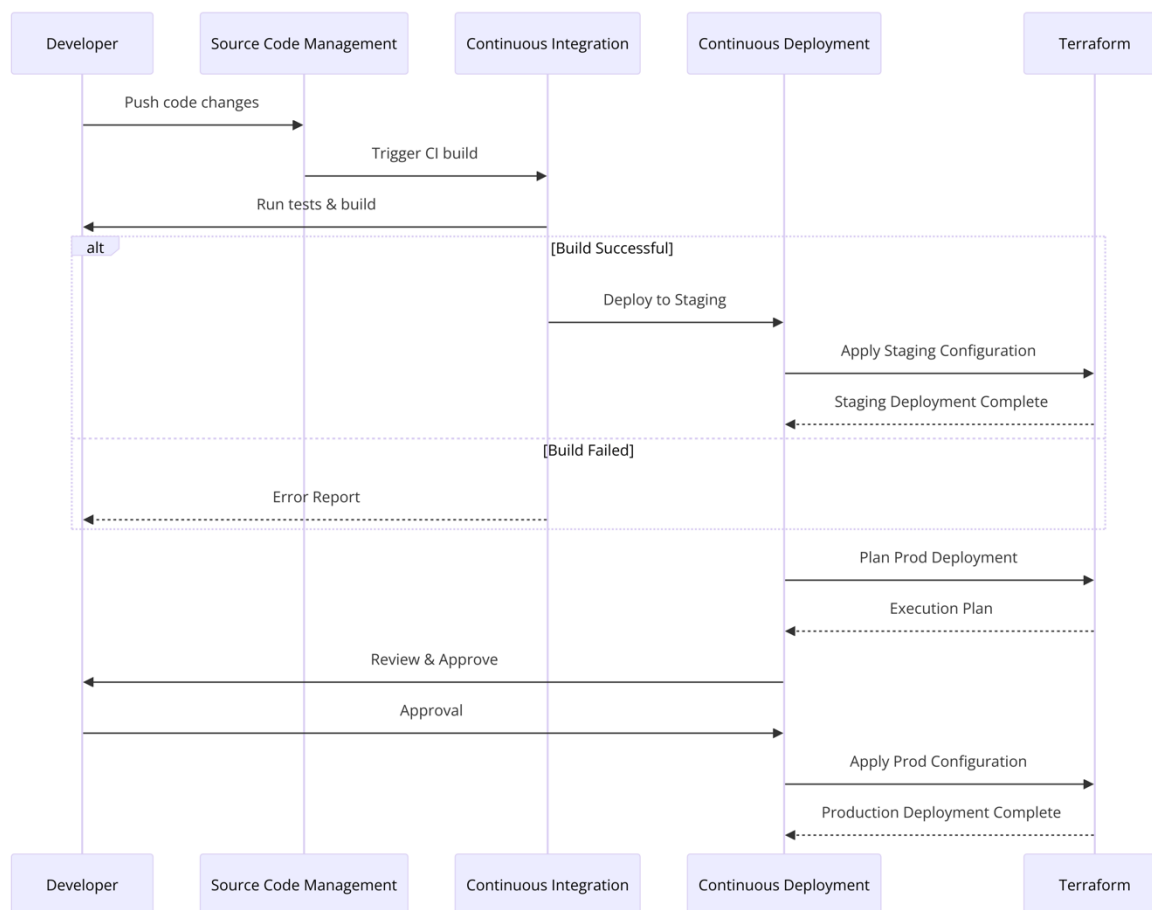


Figure 3: Illustration of a CI/CD pipeline with Terraform stages

6. Conclusion

Implementing IaC using Terraform offers significant advantages in terms of efficiency, consistency, and reliability in infrastructure management. By adhering to best practices from code organization to operation and maintenance, organizations can harness the full potential of IaC. The shift to IaC, however, is not just a technical change but also a cultural one, requiring teams to adopt new workflows and collaboration methods. As Terraform and similar tools evolve, they will continue to shape the landscape of infrastructure management, making adherence to these best practices even more crucial.

This detailed overview provides a comprehensive guide to understanding and implementing IaC with Terraform, emphasizing the importance of best practices in maximizing the benefits of this transformative approach.

References

[1] HashiCorp Official Documentation. <https://developer.hashicorp.com/terraform/install>

[2] Gustian, D., Fitriasia, Y., Novayani, W., & Purwantoro, S. E. S. G. S. (2023). Implementasi Automation Deployment pada Google Cloud Compute VM menggunakan Terraform. Implementasi Automation Deployment pada Google Cloud Compute VM menggunakan Terraform.

[3] Bailuguttu, S., Chavan, A. S., Pal, O., Sannakavalappa, K., & Chakrabarti, D. (2023). Comparing performance of bastion host on cloud using Amazon web services vs terraform. Comparing performance of bastion host on cloud using Amazon web services vs terraform.

[4] Bahaweres, R. B., & Najib, F. M. (2021). Provisioning of Disaster Recovery with Terraform and Kubernetes: A Case Study on Software Defect Prediction. Provisioning of Disaster Recovery with Terraform and Kubernetes: A Case Study on Software Defect Prediction.

[5] Hu, H., Bu, Y., Wong, K., Sood, G., Smiley, K., & Rahman, A. (2021). Characterizing Static Analysis Alerts for Terraform Manifests: An Experience Report. Characterizing Static Analysis Alerts for Terraform Manifests: An Experience Report.

- [6] Infrastructure as Code: Managing Servers in the Cloud by Kief Morris <https://www.oreilly.com/library/view/infrastructure-as-code/9781491924334/>
- [7] The Infrastructure as Code Handbook by Kief Morris <https://www.oreilly.com/library/view/the-infrastructure-as/9781492046308/>
- [8] Infrastructure as Code: A DevOps Approach by Kief Morris <https://www.oreilly.com/library/view/infrastructure-as-code/9781491924334/>
- [9] Automation Best Practices for DevOps by Red Hat <https://www.redhat.com/en/topics/devops>
- [10] Automation Best Practices for IT Operations by Puppet <https://www.puppet.com/resources/report/automation-best-practices-for-it-operations/>