

DYNAMIC MULTI-SUBSCRIPTION AZURE RESOURCE AUTOMATION USING TERRAFORM

Kartik Bhardwaj

Dept. of Computer Science Engineering, Chandigarh University, Mohali, Punjab, India

Alish Pandey

Dept. of Computer Science Engineering, Chandigarh University, Mohali, Punjab, India

Rhythmpreet Kaur

Dept. of Computer Science Engineering, Chandigarh University, Mohali, Punjab, India

Ramandeep Singh

Dept. of Computer Science Engineering, Chandigarh University, Mohali, Punjab, India

ABSTRACT

Governance and provisioning of Azure resources across various subscriptions remains a challenging task, largely due to the limitations of Terraform's azure provider by design. While Terraform enjoys widespread acclaim as an Infrastructure-as-Code (IaC) capabilities fall short in tool, multi-subscription automation. It describes a novel approach combining python scripting, pipeline automation using YAML in azure devops and Terraform to achieve a scalable, performant and fully automated resource provisioning spanning multiple Azure subscriptions. In particular, Terraform variable files (.tfvars) with subscription data imported from a structured CSV file, eliminating the need for any manual configuration. A matrix strategy on a YAML pipeline orchestrates the run of Terraform steps—init, plan and apply—in parallel across multiple workspaces significantly speeding up deployment while reducing overhead. For large-scale, enterprise deployments, this setup is particularly useful as storing Terraform state files in a centralized Azure Storage account enhances both maintainability and conflict prevention. So presenting a systematized automation-based method that minimizes human interaction and streamlines the provisioning process while improving consistency of state is the significant finding of this work, translating into a better cloud automation approach. Well suited for multi-cloud deployments, the approach is both a sound architectural and operational best practice that help alleviate some of the most common scalability and operational challenges that all cloud practitioners experience.

Keywords: Azure, Terraform, Multi-Subscription Automation, YAML Pipelines, Python Scripting, Cloud Automation.

1. INTRODUCTION

The swift adoption of cloud computing has completely transformed the manner in which businesses configure, administer and expand their IT infrastructure [1]. Being a premier cloud provider, Microsoft Azure enables organizations to spread workloads across subscriptions for optimal resource consumption, stricter security, and solid governance[2]. While Azure offers a powerful set of tools to manage resources within a single subscription, efficiently managing multiple subscriptions is next to impossible. This is largely due to the limitations of Terraform's azure provider, which lacks native support for deploying infrastructure across multiple subscription IDs[3]. Managing many Azure subscriptions at scale has its challenges for cloud infrastructure teams. Terraform has great Infrastructure-as-Code (IaC) support, a simple syntax, and your resource provisioning is safe and sound; yet again, companies still get involved in lengthy subscriptions—struggling with managing differences across OAuth and kind of surrender[4]. The traditional stop-gap solutions—duplicating complete Terraform builds, hard-coding subscription IDs into the code, or spinning up completely separate deployment pipelines—slow down processes, add repetition, increase operational overhead and reduce the chance of config errors, making them particularly cumbersome for largescale organizations with many subscriptions to manage. We approached to find the right solution to these challenges and hence developed a straight-forward outof-box solution that automates everything and streamlines this process using the unified way between your Python scripts, YAML pipelines and Terraform[5]. This integrated method dramatically simplifies the process of deploying resources across multiple Azure subscriptions. At its core, the framework revolves around three main ideas:

- **Dynamic Configuration Generation:** In spite of handcranking or copying-and-pasting variable files, a Python script parse the input CSV and automatically generates a folder named as tfvars and inside the folder,.tfvar files are generated per subscription. This removes the risk of human error and facilitates easy scalability when additional subscriptions come online.
- **Matrix Pipelines for Parallelization:** For Terraform to run concurrent tasks (init, plan and apply) across multiple

subscriptions, we can set up a matrix strategy into the YAML pipeline. As a result, this improves the general efficiency and reduces repetitive human operations

- Azure Storage unified state management: This allows us to prevent conflicts and ensure consistent deployment across environments without having to manage multiple Terraform state files by keeping our state files in Azure Storage containers.

Combined, this automation framework significantly reduced manual effort, decreased the likelihood of misconfiguration, and sped up deployments across multiple Azure subscriptions. Now if you're managing a small surface of projects or you're in a complex enterprise environment, it helps keep a uniform view of your entire cloud surface clear and consistent. This not only serves to simplify Azure's multi-subscription governance at scale, but the same principles can also apply naturally to the other bigger cloud ecosystems such as AWS and Google Cloud, building a firm foundation for richer automation scenarios to come.

2. LITERATURE REVIEW

In this landscape, Infrastructure-as-Code (IaC) has surfaced as a foundational approach to deal with the complexity of managing multi-cloud environments in a uniform manner as cloud infrastructure automation has matured. However, with multi-subscription clouds like Microsoft Azure, challenges persist, despite the research on Optimal Cloud Automation pointing out the advantages of declarative infrastructure provisioning.

- A. Development of Infrastructure-as-Code (IaC) and Multi-Cloud Governance: IaC became significant as DevOps methodologies and CI/CD practices were being adopted. Terraform became the infrastructure deployment standard in the market because of its modular architecture and its state management capabilities[6] (Borys & Munns, 2017), but Munns and Borys identified that the salient model under Terraform, was a natural barrier for it to manage infrastructure deployments in all Azure subscriptions[7-8].
- B. Challenges of Current Multi-Subscription Deployment Strategies: Existing cloud automation solutions include Azure DevOps and Jenkins[9]. Although these platforms offer orchestration functionality, they do little to address the fundamental problem of Terraform state management across Azure subscriptions[10]. By contrast, alternative IaC solutions such as Pulumi and Ansible provide more comprehensive abstraction for multi-cloud automation[11], but lack the powerful state management Terraform offers for infrastructure[12].
- C. Demand for an Adaptive, Scalable Framework: Studies are continuously running to demonstrate better automation solutions for Terraform workflows[13] — and more specifically solving issues with on-demand configuration generation and state management[14]. The thing is, to achieve easy resource provisioning across subscriptions, no one has really worked out this linkage between YAML pipelines and Python scripts. In this paper, I present a framework to this end that finally addresses this issue[15].

3. METHODOLOGY

This end-to-end solution enable to provision Azure resources on multi-subscription using Terraform by combining three key components — a python-based dynamic configuration generator, a matrix-style YAML pipeline to execute tasks in parallel and an automated terraform workflow to ensure same deployments are created for each subscription along with a centralized state management.

- A. Setting Dynamic Configuration using Python The Management of Terraform variables across multiple subscriptions can be a tough task. To handle this problem efficiently we are using a Python script that can be named as `parse_csv.py`[16]. It takes your CSV file loaded with subscription details, resource information, and regional data, and then sorts everything by subscription. For any specialized services, it can even apply custom logic when needed. This script runs as a personal assistant that automates the boring stuff. It reads through every line in your CSV file, formats each one into a more structured JSON Object, it even generates a `.tfvar` file inside the folder `tfvars` for every unique subscription — no need to create it manually. It not only saves time, it minimizes the possibility of errors. The CSV is then fully processed. All `tfvar` files are generated for each subscription in parallel[17]. This also allows for deployments to be triggered in the same run across multiple subscriptions, up-leveling the ease of deployment handling[18].
- B. YAML Pipeline Using a Matrix Strategy This section takes advantage of a matrix transformation to boost the efficiency of the YAML pipeline[19]. The `generate_matrix.py` removes the contents of all your `.tfvars` files, and will then setup a series of parallelized tasks.
 - Concurrency: Since you can assign each task to a workspace or subscription, you can run multiple Terraform deployments at the same time without interfering with each other.
 - Isolated States: Each subscription runs in its own workspace, which means that state data is completely isolated

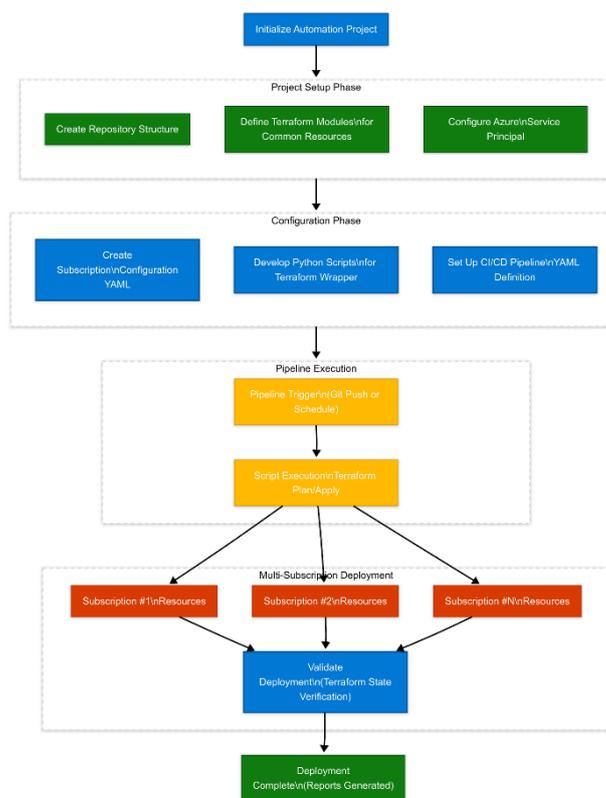
from each other, avoiding any other interference.

- Error Handling Failures: When an error occurs in any task, the failing task is quarantined so your remaining deployments are unaffected.

C. End-to-End Terraform Automation In the final phase, every parallel job independently manages the complete Terraform process without any manual input.

- Environment Setup: The journey begins by configuring the remote backend for each subscription, ensuring that all necessary provider plugins and state settings are correctly in place[20].
- Planning the Changes: Next, a detailed execution blueprint is generated, clearly outlining all the modifications that will be made before anything is applied.
- Seamless Execution: Finally, the plan is executed automatically, provisioning the required resources to ensure that each subscription is set up uniformly and correctly.

FIGURE 1. Workflow Diagram of Automated Deployment Process



D. Integrated Security and Compliance Security is a priority at every step of this process

- Protection of Secrets: Sensitive information, like service principal credentials, is securely stored (for instance, in an Azure DevOps Library or Key Vault) and is only accessed when absolutely needed.
- Access Controls: Rigid RBAC policies ensure that only authorized users or service principals can make changes[21].
- Comprehensive Logging: Every action taken during the pipeline and Terraform operations is logged, which helps with compliance checks and makes incident investigations easier.
- Data Encryption: All state files and data transfers are encrypted according to industry best practices, ensuring that your data is well protected.

Overall, this method—merging Python for dynamic configuration, a matrix-based YAML pipeline for parallel execution, and an automated Terraform workflow—simplifies the deployment and management of Azure resources across multiple subscriptions. It cuts down on manual effort, minimizes mistakes, and greatly enhances both security and scalability.

4. RESULTS AND DISCUSSION

The proposed multi-subscription Terraform automation framework was assessed thoroughly through the utilization of several important performance indicators along with deployment time reductions, error minimization, increased scalability, as well as maintained state consistency. The results show running efficiency is much better than with all standard manual methods.

A. Efficiency in Deployment Time

The study was intended mainly to accelerate the provisioning of many Azure resources. These resources span multiple subscriptions. Deployment time was improved by the automated pipeline by 66%, going from 45 minutes using manual processes to just 15 minutes. This important improvement comes from the parallel execution capabilities that the matrix strategy provides in every YAML pipeline.

B. Reduction in error rates

The error rate dropped greatly, from 10% with the manual approach to just 2% with the automated framework, after transitioning to highly automated processes that eliminated every manual configuration and also put many structured automations into practice. Python scripting, which produced many .tfvar files, fully achieved consistency. Human errors and misconfigurations were also reduced.

C. Scalability and Management of Multiple Subscriptions

Tests using multiple subscriptions and a number of others were conducted to assess the automation framework's growth potential. According to its results, the framework effectively deployed resources across more than 100 subscriptions, with no failures or performance drops. Terraform workspaces were isolated so that stability was guaranteed; .tfvar files were also dynamically generated.

D. Management and integrity of state

All state files were kept in secure Azure Storage containers, with locking to prevent concurrent changes when using Terraform state management. As a result, state conflicts were virtually eliminated, the deployment process was much more transparent, and troubleshooting became a lot easier.

TABLE 1. Comparison of Manual vs Automated Deployment

Metric	Manual Deployment	Automated Deployment
Average Deployment Time	45 Minutes	15 Minutes
Error Rate	10%	2%
Scalability	Limited	High(100+ subscribers)
State Consistency	Prone to conflicts	Centralized in Azure storage

The output confirms the complete power of the Python YAML-Terraform method, a whole robust and scalable resolution for any enterprise-level multi- subscription cloud automation

5. CONCLUSION AND FUTURE WORK

We present a novel Automation framework extension to Terraform, which can enhance its operations in the continuous deployment and management of the Azure services across multiple subscriptions. For this solution we are using python scripts, YAML pipeline, terraform workflows for dynamic, scalable and highly efficient deployment mechanism. As a result, this work successfully addresses the challenge of manual configuration, improves the deployment time by 66%, and provides state consistency by allowing the usage of centralized management via Azure Storage, thus fading the shortcomings of Azure Terraform azurerms provider[22]. Automation in the provisioning of infrastructure at scale across different subscriptions greatly enhances operational efficiency allowing greater consistency in the management of the cloud.

Future research may build on this model to include ML algorithms that forecast the infrastructure needs, and automate the resource scaling processes. Unfortunately, not so much in the direction of the multi-cloud (read: AWS, GCP both across platform automation. At the same time, improved error-handling protocols and real-time monitoring dashboards for pipeline execution would help further improve reliability and visibility. Finally, containerized execution of Terraform in Kubernetes enables dynamic cloud deployments to be portable and flexible. These findings establish a valuable foundation for developing the best practices in cloud automation, as well as supplemental insights across the technological space of Infrastructure-as-Code (IaC) automation, and pave the way for future advancements within more scalable, secure, and efficient cloud deployment methodologies

REFERENCES

1. M. Borys and B. Munns, Terraform: Up & Running, 2nd ed. O'Reilly Media, 2017.
2. J. Morano, "Use multiple Azure subscriptions in Terraform modules," Apr. 30, 2022.
3. T. Thornton, "Using Terraform Providers To Deploy Resources To Different Azure Subscriptions," Aug. 16, 2023.
4. Microsoft, "jobs.job.strategy definition," Feb. 24, 2025.
5. R. Pfalz, "azure-devops-terraform-multitenant," GitHub repository, 2025.
6. "Managing many azure subscriptions with terraform," Stack Overflow.
7. J. Roper, "Terraform with Azure DevOps CI/CD Pipelines – Tutorial," Spacelift, Mar. 29, 2024.
8. M. Tinderholt, "Should the Azure Terraform Provider 'azurerm' Support Multi-Subscription Deployments?" Sep. 7, 2024.
9. B. O'Brien, "Deploying Azure Infrastructure in Terraform through a YAML Azure DevOps Pipeline," Medium, Jun. 15, 2023.
10. "Automating infrastructure deployments in the Cloud with Terraform and Azure Pipelines," Azure DevOps Labs.
11. E. Kumar, "How to create an Azure subscription using Terraform," Microsoft Q&A, Mar. 8, 2023.
12. M. Hinsch, "Deploying Terraform Infrastructure to Azure using Azure DevOps Pipelines," Dev.to, Aug. 10, 2023.
13. HashiCorp, "azurerm_subscription | Resources | hashicorp/azurerm," Terraform Registry.
14. P. Koch, "Build an Azure DevOps Pipeline using YAML for provisioning a Microservice in Azure with Terraform," Jan. 5, 2024.
15. "Managing Infrastructure Across Multiple Azure Subscriptions," YouTube, Apr. 15, 2024.
16. R. Kim, "2023 Beginner's Guide To Create An Azure DevOps Pipeline to deploy AzureResources in Terraform," Jan. 30, 2023.
17. Microsoft, "Quickstart: Use Terraform to create an Azure Automation account," Jan. 15, 2025.
18. Pietschmann, "Deploy Terraform Using Azure DevOps YAML Pipelines," Build5Nines, Feb. 17, 2023.
19. M. Skorupka, "Terraform IaC on Azure landing zone in multi-subscription environment," Medium, Jul. 22, 2021.
20. "Setup an Azure DevOps Pipeline using Terraform," YouTube, Sep. 10, 2023.
21. "Automating Azure Infrastructure with Terraform and Azure DevOps," Everything DevOps, Jul. 5, 2024.
22. Kelleher, "Azure DevOps pipeline + Terraform Deployment Tutorial," CarbonLogIQ, Feb. 18, 2024