

HashiCorp Terraform Infrastructure as Code (IaC) Practice Test (Sample)

Study Guide



Everything you need from our exam experts!

Copyright © 2025 by Examzify - A Kaluba Technologies Inc. product.

ALL RIGHTS RESERVED.

No part of this book may be reproduced or transferred in any form or by any means, graphic, electronic, or mechanical, including photocopying, recording, web distribution, taping, or by any information storage retrieval system, without the written permission of the author.

Notice: Examzify makes every reasonable effort to obtain from reliable sources accurate, complete, and timely information about this product.

SAMPLE

Questions

- 1. What is the recommended approach to protect sensitive data identified by the security team in Terraform state files?**
 - A. Store the state in an encrypted backend**
 - B. Always store your secrets in a secrets.tfvars file**
 - C. Delete the state file every time you run Terraform**
 - D. Edit your state file to scrub out the sensitive data**
- 2. What does the 'terraform state' command do?**
 - A. It initializes the backend for storing state files**
 - B. It is used to manipulate the Terraform state directly**
 - C. It helps in downloading the latest provider versions**
 - D. It generates reports based on infrastructure changes**
- 3. What file is typically used to specify multiple environments in Terraform?**
 - A. terraform.netvars**
 - B. terraform.tfvars**
 - C. terraform.env**
 - D. terraform.config**
- 4. What language is used to write Terraform configuration files?**
 - A. JavaScript**
 - B. Python**
 - C. HashiCorp Configuration Language (HCL)**
 - D. Ruby on Rails**
- 5. Which Terraform command is used to initialize a working directory?**
 - A. terraform start**
 - B. terraform init**
 - C. terraform setup**
 - D. terraform configure**

- 6. How can you ensure consistent environments when using Terraform?**
- A. By frequently updating configurations manually**
 - B. By using the same configuration files and state management across environments**
 - C. By deploying to multiple cloud providers simultaneously**
 - D. By using separate versions of Terraform for each environment**
- 7. What type of values can variables in Terraform hold?**
- A. Only string values**
 - B. Strings, numbers, lists, maps, and booleans**
 - C. Only numeric values**
 - D. Strings and lists only**
- 8. How can you maintain idempotency in Terraform configurations?**
- A. By writing imperative scripts**
 - B. By using resource dependencies**
 - C. By writing declarative configurations**
 - D. By running 'terraform validate'**
- 9. How can you protect sensitive information in Terraform variables?**
- A. By marking variables as sensitive in the configuration**
 - B. By encrypting the configuration files**
 - C. By using version control to store variables**
 - D. By limiting access to the Terraform state file**
- 10. What is a data source in Terraform?**
- A. A read-only reference that allows you to fetch and use information from existing resources**
 - B. A method to create new resources in your configuration**
 - C. A temporary storage for sensitive information**
 - D. A tool for managing Terraform modules**

Answers

SAMPLE

1. A
2. B
3. B
4. C
5. B
6. B
7. B
8. C
9. A
10. A

SAMPLE

Explanations

SAMPLE

1. What is the recommended approach to protect sensitive data identified by the security team in Terraform state files?

- A. Store the state in an encrypted backend**
- B. Always store your secrets in a secrets.tfvars file**
- C. Delete the state file every time you run Terraform**
- D. Edit your state file to scrub out the sensitive data**

The best approach to protect sensitive data in Terraform state files is to store the state in an encrypted backend. When using an encrypted backend, all data, including sensitive information like secrets, access tokens, and passwords, is encrypted at rest. This means that even if an unauthorized individual were to gain access to the state file, they would not be able to read the sensitive information without the appropriate decryption keys. Using an encrypted backend adds a strong layer of security and aligns with best practices in Infrastructure as Code. Many backends that support Terraform, such as AWS S3 with server-side encryption, Azure Blob Storage with encryption, or HashiCorp's own Terraform Cloud, provide options for encryption that can ensure that sensitive data is not exposed. Other approaches mentioned are less effective. Storing secrets in a separate variables file that is not encrypted does not prevent exposure if the file is accessed improperly. Deleting state files after every run would be impractical and would severely hinder Terraform's ability to track resources over time. Editing state files to remove sensitive data is risky and can lead to corruption or loss of important information needed for managing infrastructure. Thus, utilizing an encrypted backend is the most robust and recommended method for protecting sensitive data in Terraform state files.

2. What does the 'terraform state' command do?

- A. It initializes the backend for storing state files**
- B. It is used to manipulate the Terraform state directly**
- C. It helps in downloading the latest provider versions**
- D. It generates reports based on infrastructure changes**

The 'terraform state' command is primarily designed for direct manipulation of the Terraform state file. This command provides users with the capability to examine and modify the current state of their managed infrastructure. This includes functions such as pulling up the current state, manipulating the resources within that state, and making changes that might not be easily achievable through regular Terraform commands. It is a crucial tool for advanced users who need to address specific issues or make fine-tuned adjustments without applying changes through standard Terraform plans and applies. The other options focus on different commands or functionalities in Terraform. Initializing the backend for storing state files is handled by the 'terraform init' command, while downloading the latest provider versions falls under updating providers through the 'terraform init -upgrade' command. Generating reports based on infrastructure changes typically involves using 'terraform plan' or 'terraform apply', rather than directly engaging with the state file. Each of these commands serves a distinct purpose within the Terraform workflow, highlighting the specialized role that the state command plays in managing and interacting with the infrastructure state directly.

3. What file is typically used to specify multiple environments in Terraform?

- A. terraform.netvars
- B. terraform.tfvars**
- C. terraform.env
- D. terraform.config

The file that is commonly used to specify multiple environments in Terraform is `terraform.tfvars`. This file allows you to define variable values that can be passed to your Terraform configuration. When managing multiple environments, it is effective to use different `tfvars` files for each environment, such as `dev.tfvars`, `prod.tfvars`, etc. This approach helps streamline the process of deploying infrastructure by ensuring that each environment can have its own specific configurations, settings, and variable values. Using `terraform.tfvars` makes it easy to manage these configurations since Terraform automatically loads the variables defined in this file when running commands. This eliminates the need for providing values explicitly on the command line or modifying configuration files directly, allowing for a cleaner separation of environment-specific settings. While other file options exist, such as `.netvars`, `.env`, and `.config`, they do not serve the same primary function of directly storing environment-specific variable values in a standard and recognized manner as `terraform.tfvars` does in Terraform usage. This is why `terraform.tfvars` is the preferred choice for managing multiple environments effectively.

4. What language is used to write Terraform configuration files?

- A. JavaScript
- B. Python
- C. HashiCorp Configuration Language (HCL)**
- D. Ruby on Rails

Terraform configuration files are written using the HashiCorp Configuration Language (HCL). HCL is designed specifically for defining infrastructure as code in a human-readable format, which makes it easier for users to write and understand their configuration files. Its syntax is declarative, enabling users to specify what they want the infrastructure to look like rather than detailing the steps to create it. This is particularly advantageous for maintaining clarity and ease of use in complex configurations. The use of HCL also allows Terraform to perform efficient parsing and validation of configuration files, aiding in error detection and helping practitioners implement best practices in infrastructure management. Furthermore, HCL's versatility extends to supporting both JSON format for configurations, but HCL is favored for its user-friendly syntax. Other programming languages like JavaScript, Python, and Ruby on Rails were not designed for this purpose, making them unsuitable for writing Terraform configuration files. Additionally, they do not provide the same level of integration and direct applicability for infrastructure as code scenarios that HCL offers.

5. Which Terraform command is used to initialize a working directory?

- A. terraform start
- B. terraform init**
- C. terraform setup
- D. terraform configure

The command used to initialize a working directory in Terraform is indeed "terraform init." This command is essential for preparing the working directory for other Terraform commands. When you run "terraform init," it performs several important tasks including downloading the necessary provider plugins specified in your Terraform configuration files, initializing the backend for state storage, and setting up the directory structure for Terraform to work with. "Terraform init" is typically the first command you should run after creating a new Terraform configuration. It ensures that the local workspace is ready for managing resources defined in your configuration files. The other options listed do not correspond to any official Terraform commands. "terraform start," "terraform setup," and "terraform configure" are not recognized commands within the Terraform ecosystem, as they do not perform any functions related to initializing the workspace or handling the project setup. Therefore, focusing on "terraform init" is critical for understanding how to properly begin working with Terraform and manage infrastructure as code effectively.

6. How can you ensure consistent environments when using Terraform?

- A. By frequently updating configurations manually
- B. By using the same configuration files and state management across environments**
- C. By deploying to multiple cloud providers simultaneously
- D. By using separate versions of Terraform for each environment

To ensure consistent environments when using Terraform, using the same configuration files and state management across environments is the most effective approach. This practice helps maintain uniformity in the infrastructure that is provisioned, as the configuration files define the specific resources, their properties, and how they relate to each other. By leveraging the same Terraform configuration, environments can be created or modified with the same settings, which minimizes discrepancies and errors that could arise from manual updates or different configurations. Consistent state management is also crucial, as it tracks the resources created by Terraform and ensures that each environment's state reflects the actual infrastructure. Proper state management across environments avoids conflicts and confusion about the current state of resources. This practice leads to a reliable and repeatable deployment process, making it easier to manage multiple environments (such as development, testing, and production) with confidence that they are all aligned with the same specifications.

7. What type of values can variables in Terraform hold?

- A. Only string values
- B. Strings, numbers, lists, maps, and booleans**
- C. Only numeric values
- D. Strings and lists only

Variables in Terraform are versatile and can hold a wide range of value types, which is crucial for creating flexible and reusable code. Specifically, variables can hold strings, numbers, lists, maps, and booleans. This variety allows Terraform configurations to adapt to different situations and architectures. - Strings are used for textual data, such as resource names or configuration settings. - Numbers can represent integer and floating-point values, which are essential for specifying resource sizes or counts. - Lists are ordered collections of values that can be of any type, enabling complex parameterizations where multiple values are handled together. - Maps are collections of key-value pairs, allowing for structured data that is easy to manage and reference throughout the configuration. - Booleans (true or false values) are critical for conditional logic within infrastructure code. Having the capability to define variables in these diverse formats makes Terraform a powerful tool for infrastructure as code, ensuring configurations are not only dynamic but also maintainable. This flexibility is essential for accommodating various resource configurations and environments.

8. How can you maintain idempotency in Terraform configurations?

- A. By writing imperative scripts
- B. By using resource dependencies
- C. By writing declarative configurations**
- D. By running 'terraform validate'

Maintaining idempotency in Terraform configurations is fundamentally achieved by writing declarative configurations. In Terraform, a declarative approach means you define the desired state of your infrastructure rather than specifying the steps to reach that state. This allows Terraform to manage changes intelligently, ensuring that if you apply the same configuration multiple times, your infrastructure ends up in the same state without unintended side effects. When you declare resources, Terraform compares the current state of the infrastructure with the desired state specified in your configuration files. If the desired state matches the actual state, no actions are taken, which reinforces the concept of idempotency. This ensures that repeated applications of the same configuration yield the same results, preventing any unwanted modifications. On the other hand, writing imperative scripts (the first choice) involves specifying exact commands that change the infrastructure imperatively, which does not inherently guarantee that repeated executions will lead to the same outcome. Resource dependencies (the second choice) help Terraform understand the relationships between resources but do not directly ensure idempotency by themselves. Running 'terraform validate' (the fourth choice) checks the configuration for syntax and validity but does not enforce idempotency in terms of resource management. Thus, utilizing declarative configurations is the key practice for achieving idempotency in

9. How can you protect sensitive information in Terraform variables?

- A. By marking variables as sensitive in the configuration**
- B. By encrypting the configuration files**
- C. By using version control to store variables**
- D. By limiting access to the Terraform state file**

Marking variables as sensitive in the Terraform configuration is a crucial practice for protecting sensitive information. When you declare a variable as sensitive, Terraform ensures that its value is not displayed in the command line output or logs, thus reducing the risk of leaking sensitive information such as passwords, API keys, or other confidential data. This feature helps manage sensitive data more securely by preventing accidental exposure during Terraform runs. It effectively informs Terraform and anyone using the code that the variable contains sensitive information. By correctly using the sensitive attribute, you maintain a level of confidentiality throughout the infrastructure provisioning process. While other methods such as encrypting configuration files, using version control, and limiting access to the Terraform state file contribute to security, they do not directly prevent sensitive data from being exposed during Terraform operations like marking a variable as sensitive does. Encrypting configuration files does offer a layer of security but doesn't inherently prevent logging of sensitive information during execution. Utilizing version control can create risks if sensitive values are committed without consideration, and limiting access to the Terraform state file is more about managing permissions than obscuring values during execution. Thus, the most effective approach within the context provided is to specifically mark variables as sensitive in the configuration.

10. What is a data source in Terraform?

- A. A read-only reference that allows you to fetch and use information from existing resources**
- B. A method to create new resources in your configuration**
- C. A temporary storage for sensitive information**
- D. A tool for managing Terraform modules**

A data source in Terraform serves as a read-only reference that enables users to fetch and utilize information from existing resources that are managed outside of Terraform or by separate configurations. This feature is essential for integrating with existing infrastructure, allowing you to reference attributes of resources that have already been created. By using data sources, you can gather information such as instance IDs, security group attributes, or any other relevant data that you need to incorporate into your configuration without having to create those resources again. This capability promotes a modular and efficient approach to building infrastructure, as it allows for the reuse of existing elements rather than duplicating efforts. It enhances collaboration within teams and ensures that your configurations are aware of the up-to-date state of existing resources. In contrast, the other choices do not accurately describe what a data source is. Creating new resources refers to resource blocks, while temporary storage for sensitive information pertains to the use of variables or Terraform's sensitive attributes. Managing Terraform modules involves structuring and organizing configuration files, which is different from the function of a data source.