Vendor: HashiCorp

Exam Code:TA-002-P

Exam Name: HashiCorp Certified: Terraform Associate

Version: Demo

QUESTION 1

Terraform works well in Windows but a Windows server is required.

A. False

B. True

Correct Answer: A

You may see this

QUESTION 2

Which of the following is not a valid string function in Terraform?

A. split

B. join

C. slice

D. chomp

Correct Answer: C

https://www.terraform.io/language/functions

QUESTION 3

You are creating a Terraform configuration which needs to make use of multiple providers, one for AWS and one for Datadog.

Which of the following provider blocks would allow you to do this?

A.

```
provider {
  "aws" {
    profile = var.aws_profile
    region  = var.aws_region
  }


  "datadog" {
    api_key = var.datadog_api_key
    app_key = var.datadog_app_key
  }
}
```

B.

```
provider "aws" {
  profile = var.aws_profile
  region  = var.aws_region
}

provider "datadog" {
  api_key = var.datadog_api_key
  app_key = var.datadog_app_key
}
```

C.

```
terraform {
 provider "aws" {
   profile = var.aws_profile
   region  = var.aws_region
 }


provider "datadog" {
 api_key = var.datadog_api_key
   app_key = var.datadog_app_key
 }
}
```

A. Option A

B. Option B

C. Option C

Correct Answer: B

https://www.terraform.io/language/providers/configuration

QUESTION 4

Which of the below options is a valid interpolation syntax for retrieving a data source?

A. ${google_storage_bucket.backend}

B. ${azurerm_resource_group.test.data}

C. ${aws_instance.web.id.data}

D. ${data.google_dns_keys.foo_dns_keys.key_signing_keys[0].ds_record}

Correct Answer: D

Data source attributes are interpolated with the general syntax data.TYPE.NAME.ATTRIBUTE. The interpolation for a resource is the same but without the data. prefix (TYPE.NAME.ATTRIBUTE).
https://www.terraform.io/docs/configuration-0-11/interpolation.html#attributes-of-a-data-source

---

QUESTION 5

How would you reference the attribute "name* of this fictitious resource in HCL?

```
resource "kubernetes_namespace" "example" {
    name = "test"
}
```

A. resource.kubrnetes_namespace>example.name

B. kubernetes_namespace.test.name

C. kubernetes_namespace.example,name

D. data kubernetes_namespace.name

E. None of the above

Correct Answer: C

https://www.terraform.io/language/expressions/references#references-to-resource-attributes

---

QUESTION 6

In the following code snippet, the block type is identified by which string?

A. "aws_instance"

B. resource

C. "db"

D. instance_type

Correct Answer: B

QUESTION 7

What is the purpose of a Terraform workspace in either open source of enterprise?

A. Workspaces allow you to manage collections of infrastructure in state files.

B. A logical separation of business units

C. A method of grouping multiple infrastructure security policies

D. Provides limited access to a cloud environment

Correct Answer: B

QUESTION 8

How does Terraform handle working with so many providers?

A. Terraform ships with all of the plugins embedded in the Terraform binary.

B. Terraform uses a plugin architecture for providers and only installs the provider plugins required by your configuration in the configuration\\\'s working directory.

C. Terraform uses a plugin architecture for providers and only installs the provider plugins required by your configuration in a shared, system-wide plugins directory.

D. Terraform allows you to select the providers you want to support during the Terraform installation process.

Correct Answer: B

Terraform is built on a plugin-based architecture. All providers and provisioners that are used in Terraform configurations are plugins, even the core types such as AWS and Heroku. Users of Terraform are able to write new plugins in order to support new functionality in Terraform.

QUESTION 9

You are writing a child Terraform module which provisions an AWS instance. You want to make use of the IP address returned in the root configuration. You name the instance resource "main". Which of these is the correct way to define the output value using HCL2?

A.

```
output "instance_ip_addr" {
  value = "${aws_instance.main.private_ip}"
}
```

B.

```
output "instance_ip_addr" {
  return aws_instance.main.private_ip
}
```

A. Option A

B. Option B

Correct Answer: A

QUESTION 10

resource "aws_s3_bucket" "example" { bucket = "my-test-s3-terraform-bucket" ...} resource "aws_iam_role" "test_role" { name = "test_role" ...}

Due to the way that the application code is written , the s3 bucket must be created before the test role is created , otherwise there will be a problem. How can you ensure that?

A. This will already be taken care of by terraform native implicit dependency. Nothing else needs to be done from your end.

B. Add explicit dependency using depends_on . This will ensure the correct order of resource creation.

C. Create 2 separate terraform config scripts , and run them one by one , 1 for s3 bucket , and another for IAM role , run the S3 bucket script first.

D. This is not possible to control in terraform . Terraform will take care of it in a native way , and create a dependency graph that is best suited for the parallel resource creation.

Correct Answer: B

Use the depends_on meta-argument to handle hidden resource dependencies that Terraform can\\'t automatically infer.

Explicitly specifying a dependency is only necessary when a resource relies on some other resource\\'s behavior but doesn\\'t access any of that resource\\'s data in its arguments.

QUESTION 11

terraform destroy is the only way to remove infrastructure.

A. True

B. False

Correct Answer: B

QUESTION 12

The Terraform CLI will print output values from a child module after running terraform apply.

A. True

B. False

Correct Answer: A