



**Free Questions for 1Z0-1109-23 by certsinside**

**Shared by Jensen on 24-05-2024**

**For More Free Questions and Preparation Resources**

**Check the Links on Last Page**

# Question 1

---

## Question Type: MultipleChoice

---

In OCI Secret management within a Vault, you have created a secret and rotated the secret one time. The current version state shows:  
Version Number Status 2 (latest) Current 1 Previous In order to rollback to version 1, what should the Administrator do?

### Options:

---

- A- Create a new secret version 3 Pending Copy the contents of Version 1 Into version 3.
- B- From the version menu, select 'Promote to current.'  
(Correct)
- C- From the version 2 latest menu, select Rollback and select version when given the option.
- D- deprecate version 2 (latest), Create a new secret version 3, create soft link for version-3 to version 1.

### Answer:

---

B

### Explanation:

---

The correct answer is: From the version menu, select 'Promote to current.' To rollback to a previous version in OCI Secret Management within a Vault, the administrator should select the desired version from the version menu and choose the option 'Promote to current.' In this scenario, the administrator wants to rollback to version 1, so they would select version 1 from the menu and promote it to the current version. This action will make version 1 the active and current version of the secret, replacing version 2. The 'Promote to current' option allows administrators to switch between different versions of a secret and make a specific version the active one.

## Question 2

---

**Question Type: MultipleChoice**

---

company uses Oracle Cloud Infrastructure (OCI) DevOps to deploy an application to their production server. They need to make some modifications to their application code and push those changes to production automatically. How can they achieve this?

### Options:

---

- A- OCI DevOps Triggers feature can be used to automate deployment.
- B- Application code can be pushed to the Resource Manager Stack for automatic deployment.
- C- Terraform code can be packaged and pushed to the OCI Code Repository to deploy the changes.
- D- Manual builds can be run from the Build Pipelines to deploy the changes.

## Answer:

---

A

## Explanation:

---

The company can use the OCI DevOps Triggers feature to automate deployment of their application code changes to the production server. Therefore, the correct answer is: OCI DevOps Triggers feature can be used to automate deployment. OCI DevOps Triggers allow for automatic builds and de-ployments based on changes to the code repository. When a new commit is pushed to the repository, the trigger can initiate a build pipeline that creates an artifact and deploys the new version of the application to the production server. Here is the link to the official documentation on using triggers in OCI DevOps to automate application deployment: Reference: <https://docs.cloud.oracle.com/en-us/iaas/devops/using/using-triggers.htm>

## Question 3

---

### Question Type: MultipleChoice

---

You are a DevOps engineer who has recently joined a new department. You have created 10 Terraform stacks using Oracle Cloud Infrastructure (OCI) Resource Manager. Each stack creates a different set of resources in OCI for your development team. What determines the cost of these Terraform stacks?

## Options:

---

- A- Resource Manager stacks are free but you are charged for the resources they create.
- B- The cost depends on the number of lines of text in your Terraform configuration files.
- C- The cost for each stack will be higher for a Pay As You Go subscription than for monthly flex billing.
- D- The cost depends on the length of time it takes to build each resource using these Terraform stacks.

## Answer:

---

A

## Explanation:

---

The correct answer is: The cost for each stack will be higher for a Pay As You Go subscription than for monthly flex billing. When it comes to the cost of using Terraform stacks created with Oracle Cloud Infrastructure (OCI) Resource Manager, it's important to note that Resource Manager stacks are free to use. However, you will be charged for the resources that are provisioned by these stacks. The cost of these resources will depend on factors such as the types of resources created, their configurations, usage duration, and the pricing model associated with your subscription. The pricing model you choose, such as Pay As You Go or monthly flex billing, will affect the cost of the resources provisioned by your Terraform stacks. Pay As You Go typically incurs usage-based charges, where you pay for the actual consumption of resources. Monthly flex billing, on the other hand, provides predictable costs based on fixed monthly commitments. The number of lines of text in your Terraform configuration files or the time it takes to build resources does not directly determine the cost. It's the actual usage and configuration of provisioned resources that impact the cost.

## Question 4

---

### Question Type: MultipleChoice

---

A DevOps Engineer is tasked with providing a solution, which will help in easy management of deployed applications and troubleshoot them on the Oracle Cloud Infrastructure Container Engine for Kubernetes (OKE). What are three actions the DevOps Engineer must perform to accomplish the given task? (Choose three.)

#### Options:

---

- A-** Manually deploy the Kubernetes dashboard on an existing cluster and access it using the URL: `dashboard:/proxy/#!/login`  
`http://localhost:8001/api/v1/namespaces/kube-system/services/https:kubernetes-dashboard:/proxy/#!/login`.
- B-** Use the default dashboard that comes configured with the Kubernetes implementation on the Oracle Cloud Infrastructure Container Engine for Kubernetes (OKE).
- C-** Create a service account and the clusterrolebinding, obtain an authentication token for the service account using `kubectl` command, and run a `kubectl proxy` command to enable the kubernetes dashboard
- D-** Automatically deploy the Kubernetes dashboard during cluster creation, create the cluster using the API and set the `isKubernetesDashboardEnabled` attribute to true and access it using the `http://localhost:8001/api/v1/namespaces/kube-system/services/https:kubernetes-dashboard:/proxy/#!/login`
- E-** Automatically deploy the Kubernetes dashboard during cluster creation, create the cluster using the API and set the `isKubernetesDashboardEnabled` attribute to true and access it using the URL: `http://localhost:8001/api/v1/namespaces/kube-dashboard/services/http:kubernetes-dashboard:/proxy/#!/login`

**F-** Manually deploy the Kubernetes dashboard on an existing cluster and access it using the URL:  
hs://localhost:8001/api/vi/namespaces/kube-dashboard/services/httparkubernetedashboard:/proxy/#!/login

### **Answer:**

---

C, D, F

### **Explanation:**

---

The three actions that the DevOps Engineer must perform to easily manage and troubleshoot applications on Oracle Cloud Infrastructure Container Engine for Kubernetes (OKE) are: Create a service account and the clusterrolebinding, obtain an authentication token for the service account using the kubectl command, and run a kubectl proxy command to enable the Kubernetes dashboard. This allows for easy access to the dashboard and management of deployed applications. Automatically deploy the Kubernetes dashboard during cluster creation, create the cluster using the API, and set the iskubernetesDashboardEnabled attribute to true. This ensures that the Kubernetes dashboard is automatically deployed and accessible. Manually deploy the Kubernetes dashboard on an existing cluster and access it using the appropriate URL. This involves deploying the dashboard manually and accessing it through the specified URL, which allows for management and troubleshooting of applications. Using these actions, the DevOps Engineer can effectively manage and troubleshoot applications deployed on OKE, leveraging the Kubernetes dashboard for enhanced visibility and control.

## **Question 5**

---

**Question Type: MultipleChoice**

---

You host your application on a stack in Oracle Cloud Infrastructure (OCI) Resource Manager. Due to recent growth in your user base, you decide to add a CIDR block to your VCN, add a subnet, and provision a compute instance in it. Which statement is true?

### Options:

---

- A-** You need to provision a new stack because Terraform uses immutable infrastructure.
- B-** You can provision the new resources in the OCI console and add them to the stack with Drift Detection.
- C-** You cannot provision the new resources in the OCI console first, then later add them to the Terraform configuration and state.
- D-** You can make the changes to the Terraform code, run an Apply job, and Resource Manager will provision the new resources.

### Answer:

---

A

### Explanation:

---

The correct statement is: You need to provision a new stack because Terraform uses immutable in-frastructure. In Oracle Cloud Infrastructure (OCI) Resource Manager, Terraform uses the concept of immutable infrastructure, which means that any changes to the infrastructure are managed through the Terraform code. In this scenario, if you want to add a CIDR block, subnet, and compute instance to your VCN, you would need to make the necessary changes to your Terraform code, create a new stack in Resource Manager, and deploy the updated code. This ensures that the infra-structure is created consistently and according to the desired state defined in the Terraform code. Simply provisioning the new resources in the OCI console and later adding them to the Terraform configuration and



state would not be the recommended approach in this case.

## Question 6

---

**Question Type:** MultipleChoice

---

A development team leveraging the Oracle Cloud Infrastructure DevOps service is having trouble getting their build pipeline to complete successfully. Which two situations might be the problem? (Choose two.)

### Options:

---

- A-** Their source code and Kubernetes manifest are in different Git repositories.
- B-** The build spec.yaml file is in the root directory of their Git repository, and they didn't specify a path to it.
- C-** They forgot to export a required variable in the build\_spec.yaml file.
- D-** Their build specification file is available in a different directory of their Git repository, and there is no reference to its location.
- E-** They did not export a vault variable in the vaultVariables section of the build\_spec.yaml file.

### Answer:

---

D, E

### **Explanation:**

---

The two situations that might be causing the problem with the build pipeline in the Oracle Cloud Infrastructure DevOps service are: They did not export a vault variable in the vaultVariables section of the build\_spec.yaml file. When using vault variables in the build specification file (build\_spec.yaml), it is necessary to export the vault variables in the vaultVariables section of the file. If this step is missed, the pipeline may fail due to missing or inaccessible vault variables. Their build specification file is available in a different directory of their Git repository, and there is no reference to its location. The build specification file (build\_spec.yaml) should be properly referenced in the pipeline configuration. If the file is located in a different directory than the default location or if its location is not specified correctly in the pipeline configuration, the pipeline may fail to find and execute the build specification, leading to a failure. To resolve these issues, the development team should ensure that they export the required vault variables in the build specification file and correctly reference the location of the build specification file in the pipeline configuration.

## **Question 7**

---

**Question Type:** MultipleChoice

---

Your customer has deployed their microservices based application on Oracle Container Engine for Kubernetes (OKE) and they are using Oracle Cloud Infrastructure Registry (OCIR) service as their Docker image repository. They have deployed the OKE cluster using the 'custom create' option, and their Virtual Cloud Network (VCN) has three public subnets with associated route tables, security lists, and an internet gateway. They are facing an issue where their application containers are failing to deploy. Upon investigation, they learn that

the images are not getting pulled from the designated OCIR repository. The YAML configuration has the correct path to the images. What is a valid concern that needs to be further investigated?

### Options:

---

- A-** The OKE cluster needs to have a secret with credentials of their OCIR repository and use that secret in the Kubernetes deployment manifest.
- B-** They need to add IAM credentials for each user that deploys applications to the OKE cluster.
- C-** The VCN hosting the OKE cluster worker nodes needs to have a NAT gateway to access OCIR repositories.
- D-** They need to add a security list rule for TCP port 22 to connect to the OCIR service.

### Answer:

---

A

### Explanation:

---

A valid concern that needs to be further investigated in this scenario is whether the OKE cluster has a secret with credentials of the Oracle Cloud Infrastructure Registry (OCIR) repository and if that secret is being used in the Kubernetes deployment manifest. When deploying an application on OKE and pulling images from OCIR, the cluster needs to authenticate and authorize access to the OCIR repository. This is typically done by creating a Kubernetes secret that contains the credentials (authentication token or username/password) required to access the repository. The secret is then referenced in the Kubernetes deployment manifest to allow the cluster to pull the images. If the images are not getting pulled from the designated OCIR repository, it suggests that the OKE cluster

might be missing the necessary secret with the OCIR credentials or the secret is not properly referenced in the deployment manifest. Further investigation should focus on ensuring the existence and correct configuration of the secret and its usage in the deployment process.

## Question 8

---

**Question Type:** MultipleChoice

---

(CHK\_1>3) What cannot be specified in a Schema Document for Oracle Cloud Infrastructure (OCI) Resource Manager?

### Options:

---

- A- dependency relationships between variables.
- B- a logo for the Resource Manager template
- C- pattern validations for string-type variables
- D- information about the application such as its name, description, and version.
- E- permissions for which OCI users may provision the template

## Answer:

---

A

## Explanation:

---

The correct answer is: dependency relationships between variables. In an Oracle Cloud Infrastructure (OCI) Resource Manager Schema Document, you can specify various aspects of your template, such as information about the application, permissions, logo, and pattern validations for string-type variables. However, dependency relationships between variables cannot be specified in a Schema Document. Dependency relationships between variables are typically defined in the Terraform configuration files themselves rather than in the Schema Document. Terraform allows you to express dependencies between resources and variables directly within the configuration using features like interpolation and variable references. The Schema Document in OCI Resource Manager primarily focuses on providing metadata and validation rules for the template inputs, but it does not include features for defining dependencies between variables.

## Question 9

---

### Question Type: MultipleChoice

---

Which TWO components are optional while creating Monitoring Query Language (MQL) expressions in Oracle Cloud Infrastructure Monitoring service? (Choose two.)

## Options:

---

- A- Dimensions
- B- Metric
- C- Grouping Function
- D- Statistic
- E- Interval

## Answer:

---

A, C

## Explanation:

---

When creating Monitoring Query Language (MQL) expressions in Oracle Cloud Infrastructure Monitoring service, the optional components are: Dimensions: Dimensions provide additional context or filters for the metrics being queried. They allow you to narrow down the scope of the query by specifying specific resources, regions, or other properties. Grouping Function: The grouping function is used to aggregate or group the data based on specified dimensions. It allows you to perform calculations or analysis on a subset of data and present the results in a summarized form. The components that are not optional when creating MQL expressions are: Statistic: The statistic component is mandatory and represents the specific metric or data point you want to retrieve or analyze. It can be a simple statistic like average, sum, count, etc., or a complex expression involving mathematical or logical operations. Metric: The metric component is also mandatory and refers to the specific metric you want to monitor or analyze. It represents the data being collected and reported by the monitoring service, such as CPU utilization, network traffic, or custom metrics. Interval is not a component of MQL expressions. It refers to the time range or period over which the query is executed and is not specified within the MQL expression itself.

Reference: <https://docs.oracle.com/en-us/iaas/Content/Monitoring/Reference/mql.htm>

**To Get Premium Files for 1Z0-1109-23 Visit**

**<https://www.p2pexams.com/products/1z0-1109-23>**

**For More Free Questions Visit**

**<https://www.p2pexams.com/oracle/pdf/1z0-1109-23>**

