



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

**Shared by Dillon on 07-08-2023**

**For More Free Questions and Preparation Resources**

**Check the Links on Last Page**

# Question 1

---

**Question Type:** MultipleChoice

---

As a cloud engineer, you are responsible for managing a Kubernetes cluster on the Oracle Cloud Infrastructure (OCI) platform for your organization. You are looking for ways to ensure reliable operations of Kubernetes at scale while minimizing the operational overhead of managing the worker node infrastructure. Which cluster option is the best fit for your requirement?

## Options:

---

- A-** Using OCI OKE virtual nodes to eliminate worker node infrastructure management.
- B-** Using OCI OKE managed nodes with cluster autoscalers to eliminate worker node infrastructure management.
- C-** Using Kubernetes cluster add-ons to automate worker node management.
- D-** Creating and managing worker nodes using OCI compute instances.

## Answer:

---

A

## Explanation:

---

The cluster option that is the best fit for ensuring reliable operations of Kubernetes at scale while minimizing the operational overhead of managing the worker node infrastructure is using OCI OKE virtual nodes to eliminate worker node infrastructure management. OKE is a fully managed service that allows you to run and manage your Kubernetes cluster on OCI. A Kubernetes cluster consists of a control plane and a set of worker nodes that run your containerized applications. OKE provides a highly available and secure control plane that is managed by Oracle, while you are responsible for managing the worker node infrastructure. However, OKE also supports virtual nodes, which are serverless compute resources that are automatically provisioned and scaled by OCI based on your application workload demands. Virtual nodes eliminate the need for managing worker node infrastructure, such as security patches, updates, scaling, etc. Virtual nodes also offer a serverless Kubernetes experience, where you can focus on developing and deploying your applications without worrying about the underlying infrastructure. Verified Reference: [Container Engine for Kubernetes - Oracle Cloud Infrastructure Developer Tools], [Virtual Nodes - Oracle Cloud Infrastructure Container Engine for Kubernetes]

## Question 2

---

**Question Type:** MultipleChoice

---

As a small company that wants to adopt a DevOps framework and a consumption-based pricing model, which Oracle Cloud Infrastructure service can be used as a target deployment environment, providing features like automated rollouts and rollbacks, self-healing of failed containers, and configuration management, without the overhead of managing security patches and scaling?

**Options:**

---

- A- OCI Container Engine for Kubernetes (OKE) with managed nodes
- B- Compute Instance Group
- C- OCI Container Instances
- D- OCI Serverless Functions
- E- OCI Container Engine for Kubernetes (OKE) with virtual nodes

### Answer:

---

E

### Explanation:

---

The OCI service that can be used as a target deployment environment for adopting a DevOps framework and a consumption-based pricing model, while providing features like automated rollouts and rollbacks, self-healing of failed containers, and configuration management, without the overhead of managing security patches and scaling, is OCI Container Engine for Kubernetes (OKE) with virtual nodes. OKE is a fully managed service that allows you to run and manage your containerized applications on OCI using Kubernetes, an open-source system for automating deployment, scaling, and management of containerized applications. OKE provides features such as automated rollouts and rollbacks, self-healing of failed containers, configuration management, service discovery, load balancing, etc. OKE also supports virtual nodes, which are serverless compute resources that are automatically provisioned and scaled by OCI based on your application workload demands. Virtual nodes eliminate the need for managing worker node infrastructure, such as security patches, updates, scaling, etc. Virtual nodes also offer a consumption-based pricing model, where you only pay for the resources you consume when your containers are running. Verified Reference: [Container Engine for Kubernetes - Oracle Cloud Infrastructure Developer Tools], [Virtual Nodes - Oracle Cloud Infrastructure Container Engine for Kubernetes]

## Question 3

---

**Question Type:** MultipleChoice

---

A fresher joined a company who made a mistake while doing yaria: to build\_spec.yaml file. As a consequence, build pipelines started failing. What is the root cause for this error committed by the fresher? (Choose the best answer.)

### Options:

---

- A- No errors. False alarm by OCI monitoring service
- B- Corrupt build\_spec.yaml
- C- Multiple comment line in build\_spec.yaml file
- D- Expected ion-exported variable of a build to be persistent throughout multiple pipeline

### Answer:

---

D

### Explanation:

---

The root cause for the error committed by the fresher is that the expected input/exported variable of a build is not persistent throughout multiple pipelines. This means that the value set for a variable in one pipeline is not carried over to subsequent pipelines, leading to failures in the build pipelines.

## Question 4

---

### Question Type: MultipleChoice

---

You're working with Docker images and containers, and you need to get rid of an image that's no longer needed. However, when you use the command to remove the image, you get the following error: `# docker rmi d1230455e12 Error response from daemon: conflict: unable to delete d1230455e12 must be forced image is being used by stopped container of1262bd1285` Which steps would you take to delete the specified image?

### Options:

---

- A-** To determine which container is using the image, use the `"docker ps -a"` command. Stop the running container with `'docker stop 0f1262bd1285'`, delete it with `'docker rm 0f1262bd1285'`, and delete the docker image with `'docker rmi d123f4e55e12'`.
- B-** To determine which container is using the image, use the `"docker ps a"` command Kill the running container with `'docker Kill of 1262bd1205'`, delete it with `'docker rmc 0r1262bd1285'`, and delete the docker image with `'docker mi d123f4e55e12'`.

**C-** To determine which container is using the image, use the "docker ps -a" command. Stop the running container with the command 'docker stop d123f4e55e12', delete it with 'docker rm 0f1262bd1285', and delete the docker image with the command 'docker rmi 0f1262bd1285'.

**D-** To determine which container is using the image, use the "docker ps -a" command. To delete the container, use the 'docker rm 01262bd1285' command and to delete the docker image, use the 'docker rmi d12314e55e12' command.

### Answer:

---

D

### Explanation:

---

The steps that you would take to delete the specified image are:

To determine which container is using the image, use the "docker ps -a" command. This command will list all the containers in your system, including their ID, image name, status, etc. You can look for the container that has the image name d123f4e55e12 in its output.

To delete the container, use the 'docker rm 0f1262bd1285' command. This command will remove the container with the ID 0f1262bd1285 from your system. You can also use the -f flag to force remove the container if it is running or paused.

To delete the docker image, use the 'docker rmi d123f4e55e12' command. This command will remove the image with the ID d123f4e55e12 from your system. You can also use the -f flag to force remove the image if it is being used by another container. Verified Reference: [docker ps - Docker Documentation], [docker rm - Docker Documentation], [docker rmi - Docker Documentation]

## Question 5

---

### Question Type: MultipleChoice

---

As a developer working on a software project to be deployed on OKE, you have created a Helm chart for your application. You want to deploy the chart using OCI DevOps service. Which statement is true about deploying a Helm Chart using OCI DevOps service?

#### Options:

---

- A- The values.yaml file is a generic file located in the OCI Container Registry. You must create a reference to this file.
- B- The Helm charts must be located in the OCI Container Registry repository for deployment.
- C- You cannot package and push the Helm chart to the Container Registry from the build pipeline.
- D- The Helm Upgrade options cannot be used to change the configuration of the Helm re-lease.

#### Answer:

---

B

#### Explanation:

---

The statement that is false about deployment pipeline in OCI DevOps is that using deployment pipeline, you can deploy artifacts to Kubernetes cluster, Instance Group, and OCI Compute Instances. This statement is false because using deployment pipeline, you can



deploy artifacts to Kubernetes cluster and Instance Group, but not to OCI Compute Instances. A deployment pipeline is a component of OCI DevOps service that allows you to automate the deployment of your artifacts to various target environments. A target environment is a destination where you want to deploy your artifact, such as a Kubernetes cluster or an Instance Group. An Instance Group is a group of compute instances that share the same configuration and are managed as a single unit. However, OCI DevOps service does not support deploying artifacts directly to individual OCI Compute Instances. Verified Reference: [Deployment Pipelines - Oracle Cloud Infrastructure DevOps], [Target Environments - Oracle Cloud Infrastructure DevOps], [Instance Groups - Oracle Cloud Infrastructure DevOps]

## Question 6

---

**Question Type:** MultipleChoice

---

You are part of the DevOps team and troubleshooting an issue related to a newly deployed web application. The clients for the web application have reported failures with creating records into the application over an HTTPS connection. The current logs collected by the Oracle Cloud Infrastructure (OCI) Logging service is not providing much information related to the issue. You have been asked to enable specific logs applicable to services along with configuring an alarm to monitor any new failures. Which two steps can you perform to meet this requirement?

**Options:**

---

**A-** Create a Rule Condition for DevOps Events

(com.oraclecloud.devopsbuild.createconnection.begin,com.oraclecloud.devopsbuild.createconnection.end) and enable Notifications Service.

**B-** Install the OCI compute agent software on client systems, enable Custom log and create an agent configuration selecting log path.

**C-** Create custom filters with required data fields (for example: source, time, statusCode, message) to filter log messages, configure Service Connector with Monitoring for creating an Alarm.

**D-** Enable Custom log specific to the application server log and create an agent configuration with selecting log paths.

**Answer:**

---

B, C

**Explanation:**

---

The steps that you can perform to enable specific logs applicable to services along with configuring an alarm to monitor any new failures are:

Install the OCI compute agent software on client systems, enable Custom log and create an agent configuration selecting log path. The OCI compute agent is a software component that runs on your compute instances and collects logs from various sources, such as files, syslog, Windows Event Log, etc. You can use the OCI compute agent to enable Custom log, which is a type of log that allows you to define your own log source and format. You can also create an agent configuration that specifies the log path, log group, and log name for your Custom log.

Create custom filters with required data fields (for example: source, time, statusCode, message) to filter log messages, configure Service Connector with Monitoring for creating an Alarm. A custom filter is a query that allows you to filter and analyze your log messages based

on various data fields, such as source, time, level, message, etc. You can use custom filters to search for specific patterns or conditions in your logs, such as failures or errors. You can also configure a Service Connector with Monitoring, which is a component that allows you to transfer data from one OCI service to another. You can use a Service Connector with Monitoring to send your filtered log messages to the OCI Monitoring service, which is a service that allows you to create metrics and alarms based on your logs. You can then create an Alarm, which is a rule that triggers an action when a metric meets a specified threshold. Verified Reference: [Compute Agent - Oracle Cloud Infrastructure Logging], [Custom Logs - Oracle Cloud Infrastructure Logging], [Custom Filters - Oracle Cloud Infrastructure Logging], [Service Connectors - Oracle Cloud Infrastructure Logging], [Monitoring - Oracle Cloud Infrastructure Logging], [Alarms - Oracle Cloud Infrastructure Logging]

## Question 7

---

**Question Type:** MultipleChoice

---

You host a microservices based application on the Oracle Cloud Infrastructure Container Engine for Kubernetes (OKE). Due to increased popularity of your application, you need to provision more resources to meet the growing demand. Which three statements are true for the given scenario?

### Options:

---

**A-** Enable autoscaling by autoscaling Pods by deploying Kubernetes Autoscaler to collect resource metrics from each worker node in the cluster.

- B-** Enable cluster autoscaling by autoscaling node pools by deploying the Kubernetes Autoscaler to automatically resize a cluster's node pools based on application workload demands.
- C-** Scale a cluster up and down by changing the number of node pools in the cluster.
- D-** Enable cluster autoscaling by autoscaling node pools by deploying Kubernetes Metrics Server and using the Kubernetes Vertical Pod Autoscaler to adjust the resource re-quests and limits.
- E-** Scale a node pool up and down to change the number of worker nodes in the node pool, and the availability domains and subnets in which to place them.

### **Answer:**

---

A, B, E

### **Explanation:**

---

The statements that are true for scaling an OKE cluster to meet growing demand are:

Enable autoscaling by autoscaling Pods by deploying Kubernetes Autoscaler to collect resource metrics from each worker node in the cluster. Pod autoscaling is a feature that allows you to adjust the number of pods in a deployment or replica set based on the CPU or memory utilization of the pods. You can use Kubernetes Autoscaler, which is an add-on component that you can install on your OKE cluster, to collect resource metrics from each worker node and scale the pods up or down accordingly.

Enable cluster autoscaling by autoscaling node pools by deploying the Kubernetes Autoscaler to automatically resize a cluster's node pools based on application workload demands. Cluster autoscaling is a feature that allows you to adjust the number of nodes in a node pool based on the pod requests and limits of the pods running on the nodes. You can use Kubernetes Autoscaler, which is an add-on

component that you can install on your OKE cluster, to monitor the pod requests and limits and scale the node pools up or down accordingly.

Scale a node pool up and down to change the number of worker nodes in the node pool, and the availability domains and subnets in which to place them. A node pool is a group of worker nodes within an OKE cluster that share the same configuration, such as shape, image, subnet, etc. You can use OCI Console, CLI, or API to scale a node pool up and down by adding or removing worker nodes from it. You can also change the availability domains and subnets for your node pool to distribute your nodes across different fault domains. Scaling a node pool allows you to adjust your cluster capacity according to your application workload demands. Verified Reference: [Scaling Clusters - Oracle Cloud Infrastructure Container Engine for Kubernetes], [Scaling Node Pools - Oracle Cloud Infrastructure Container Engine for Kubernetes]

## Question 8

---

**Question Type:** MultipleChoice

---

You're working with Docker images and containers, and you need to get rid of an image that's no longer needed. However, when you use the command to remove the image, you get the following error: `# docker rmi d1230455e12 Error response from daemon: conflict: unable to delete d1230455e12 must be forced image is being used by stopped container of1262bd1285` Which steps would you take to delete the specified image?

## Options:

---

- A-** To determine which container is using the image, use the "docker ps -a" command. Stop the running container with 'docker stop 0f1262bd1285', delete it with 'docker rm 0f1262bd1285', and delete the docker image with 'docker rmi d123f4e55e12'.
- B-** To determine which container is using the image, use the "docker ps a" command Kill the running container with 'docker Kill of 1262bd1205, delete it with 'docker rmc 0r1262bd1285', and delete the docker image with 'docker mi d123f4e55e12'.
- C-** To determine which container is using the image, use the "docker ps -a" command. Stop the running container with the command 'docker stop d123f4e55e12', delete it with 'docker rm Of1262bd1285', and delete the docker image with the command 'dock-er rmi 0f1262bd1285'.
- D-** To determine which container is using the image, use the "docker ps -a" command. To delete the container, use the 'docker rm 011262bd1285' command and to delete the docker image, use the 'docker rmi d12314e55e12' command.

## Answer:

---

D

## Explanation:

---

The steps that you would take to delete the specified image are:

To determine which container is using the image, use the "docker ps -a" command. This command will list all the containers in your system, including their ID, image name, status, etc. You can look for the container that has the image name d123f4e55e12 in its output.

To delete the container, use the 'docker rm 0f1262bd1285' command. This command will remove the container with the ID 0f1262bd1285 from your system. You can also use the -f flag to force remove the container if it is running or paused.

To delete the docker image, use the 'docker rmi d123f4e55e12' command. This command will remove the image with the ID d123f4e55e12 from your system. You can also use the -f flag to force remove the image if it is being used by another container. Verified Reference: [docker ps - Docker Documentation], [docker rm - Docker Documentation], [docker rmi - Docker Documentation]

## Question 9

---

### Question Type: MultipleChoice

---

As a cloud engineer, you are responsible for managing a Kubernetes cluster on the Oracle Cloud Infrastructure (OCI) platform for your organization. You are looking for ways to ensure reliable operations of Kubernetes at scale while minimizing the operational overhead of managing the worker node infrastructure. Which cluster option is the best fit for your requirement?

### Options:

---

- A- Using OCI OKE virtual nodes to eliminate worker node infrastructure management.
- B- Using OCI OKE managed nodes with cluster autoscalers to eliminate worker node infrastructure management.

- C- Using Kubernetes cluster add-ons to automate worker node management.
- D- Creating and managing worker nodes using OCI compute instances.

**Answer:**

---

A

**Explanation:**

---

The cluster option that is the best fit for ensuring reliable operations of Kubernetes at scale while minimizing the operational overhead of managing the worker node infrastructure is using OCI OKE virtual nodes to eliminate worker node infrastructure management. OKE is a fully managed service that allows you to run and manage your Kubernetes cluster on OCI. A Kubernetes cluster consists of a control plane and a set of worker nodes that run your containerized applications. OKE provides a highly available and secure control plane that is managed by Oracle, while you are responsible for managing the worker node infrastructure. However, OKE also supports virtual nodes, which are serverless compute resources that are automatically provisioned and scaled by OCI based on your application workload demands. Virtual nodes eliminate the need for managing worker node infrastructure, such as security patches, updates, scaling, etc. Virtual nodes also offer a serverless Kubernetes experience, where you can focus on developing and deploying your applications without worrying about the underlying infrastructure. Verified Reference: [Container Engine for Kubernetes - Oracle Cloud Infrastructure Developer Tools], [Virtual Nodes - Oracle Cloud Infrastructure Container Engine for Kubernetes]

## Question 10

---



**Question Type: MultipleChoice**

---

You are part of the DevOps team and troubleshooting an issue related to a newly deployed web application. The clients for the web application have reported failures with creating records into the application over an HTTPS connection. The current logs collected by the Oracle Cloud Infrastructure (OCI) Logging service is not providing much information related to the issue. You have been asked to enable specific logs applicable to services along with configuring an alarm to monitor any new failures. Which two steps can you perform to meet this requirement?

**Options:**

---

- A-** Create a Rule Condition for DevOps Events  
(com.oraclecloud.devopsbuild.createconnection.begin,com.oraclecloud.devopsbuild.createconnection.end) and enable Notifications Service.
- B-** Install the OCI compute agent software on client systems, enable Custom log and create an agent configuration selecting log path.
- C-** Create custom filters with required data fields (for example: source, time, statusCode, message) to filter log messages, configure Service Connector with Monitoring for creating an Alarm.
- D-** Enable Custom log specific to the application server log and create an agent configuration with selecting log paths.

**Answer:**

---

B, C

## **Explanation:**

---

The steps that you can perform to enable specific logs applicable to services along with configuring an alarm to monitor any new failures are:

Install the OCI compute agent software on client systems, enable Custom log and create an agent configuration selecting log path. The OCI compute agent is a software component that runs on your compute instances and collects logs from various sources, such as files, syslog, Windows Event Log, etc. You can use the OCI compute agent to enable Custom log, which is a type of log that allows you to define your own log source and format. You can also create an agent configuration that specifies the log path, log group, and log name for your Custom log.

Create custom filters with required data fields (for example: source, time, statusCode, message) to filter log messages, configure Service Connector with Monitoring for creating an Alarm. A custom filter is a query that allows you to filter and analyze your log messages based on various data fields, such as source, time, level, message, etc. You can use custom filters to search for specific patterns or conditions in your logs, such as failures or errors. You can also configure a Service Connector with Monitoring, which is a component that allows you to transfer data from one OCI service to another. You can use a Service Connector with Monitoring to send your filtered log messages to the OCI Monitoring service, which is a service that allows you to create metrics and alarms based on your logs. You can then create an Alarm, which is a rule that triggers an action when a metric meets a specified threshold. Verified Reference: [Compute Agent - Oracle Cloud Infrastructure Logging], [Custom Logs - Oracle Cloud Infrastructure Logging], [Custom Filters - Oracle Cloud Infrastructure Logging], [Service Connectors - Oracle Cloud Infrastructure Logging], [Monitoring - Oracle Cloud Infrastructure Logging], [Alarms - Oracle Cloud Infrastructure Logging]

## **Question 11**

---

**Question Type: MultipleChoice**

---

As a small company that wants to adopt a DevOps framework and a consumption-based pricing model, which Oracle Cloud Infrastructure service can be used as a target deployment environment, providing features like automated rollouts and rollbacks, self-healing of failed containers, and configuration management, without the overhead of managing security patches and scaling?

**Options:**

---

- A- OCI Container Engine for Kubernetes (OKE) with managed nodes
- B- Compute Instance Group
- C- OCI Container Instances
- D- OCI Serverless Functions
- E- OCI Container Engine for Kubernetes (OKE) with virtual nodes

**Answer:**

---

E

**Explanation:**

---

The OCI service that can be used as a target deployment environment for adopting a DevOps framework and a consumption-based pricing model, while providing features like automated rollouts and rollbacks, self-healing of failed containers, and configuration

management, without the overhead of managing security patches and scaling, is OCI Container Engine for Kubernetes (OKE) with virtual nodes. OKE is a fully managed service that allows you to run and manage your containerized applications on OCI using Kubernetes, an open-source system for automating deployment, scaling, and management of containerized applications. OKE provides features such as automated rollouts and rollbacks, self-healing of failed containers, configuration management, service discovery, load balancing, etc. OKE also supports virtual nodes, which are serverless compute resources that are automatically provisioned and scaled by OCI based on your application workload demands. Virtual nodes eliminate the need for managing worker node infrastructure, such as security patches, updates, scaling, etc. Virtual nodes also offer a consumption-based pricing model, where you only pay for the resources you consume when your containers are running. Verified Reference: [Container Engine for Kubernetes - Oracle Cloud Infrastructure Developer Tools], [Virtual Nodes - Oracle Cloud Infrastructure Container Engine for Kubernetes]

## Question 12

---

**Question Type:** MultipleChoice

---

As a developer working on a software project to be deployed on OKE, you have created a Helm chart for your application. You want to deploy the chart using OCI DevOps service. Which statement is true about deploying a Helm Chart using OCI DevOps service?

**Options:**

---

**A-** The values.yaml file is a generic file located in the OCI Container Registry. You must create a reference to this file.

- B-** The Helm charts must be located in the OCI Container Registry repository for deployment.
- C-** You cannot package and push the Helm chart to the Container Registry from the build pipeline.
- D-** The Helm Upgrade options cannot be used to change the configuration of the Helm re-lease.

### **Answer:**

---

B

### **Explanation:**

---

The statement that is false about deployment pipeline in OCI DevOps is that using deployment pipeline, you can deploy artifacts to Kubernetes cluster, Instance Group, and OCI Compute Instances. This statement is false because using deployment pipeline, you can deploy artifacts to Kubernetes cluster and Instance Group, but not to OCI Compute Instances. A deployment pipeline is a component of OCI DevOps service that allows you to automate the deployment of your artifacts to various target environments. A target environment is a destination where you want to deploy your artifact, such as a Kubernetes cluster or an Instance Group. An Instance Group is a group of compute instances that share the same configuration and are managed as a single unit. However, OCI DevOps service does not support deploying artifacts directly to individual OCI Compute Instances. Verified Reference: [Deployment Pipelines - Oracle Cloud Infrastructure DevOps], [Target Environments - Oracle Cloud Infrastructure DevOps], [Instance Groups - Oracle Cloud Infrastructure DevOps]

**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>**

