



Free Questions for 1Z0-1084-23 by certsdeals

Shared by Stokes on 24-05-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Which THREE are valid statements regarding the OCI Container Engine for Kubernetes (OKE) service? (Choose three.)

Options:

- A-** You must have access to an Oracle Cloud Infrastructure tenancy. Your tenancy must have sufficient quota on different types of resources.
- B-** OKE cannot use existing network resources for the creation of a new cluster.
- C-** OKE automatically creates and configures new network resources for the new cluster.
- D-** There is a limit of three clusters within each region, but there is no limit on the number of nodes and pods you can create within each cluster.

Answer:

A, C, D

Explanation:

The valid statements regarding the OCI Container Engine for Kubernetes (OKE) service are: OKE automatically creates and configures new network resources for the new cluster. When creating a new OKE cluster, the service automatically provisions and configures the necessary network resources, such as VCNs, subnets, route tables, security lists, and load balancers, to support the cluster. Your tenancy must have sufficient quota on different types of resources. Before creating an OKE cluster, you need to ensure that your Oracle Cloud Infrastructure (OCI) tenancy has sufficient quota for the required resources, such as compute instances, block storage, networking resources, and load balancers. You must have access to an Oracle Cloud Infrastructure tenancy. To use the OKE service, you need to have access to an OCI tenancy. This means you must have a valid OCI account and the necessary permissions to create and manage resources within the tenancy. The following statements are not valid: OKE cannot use existing network resources for the creation of a new cluster. OKE creates new network resources specifically for the cluster, and it does not support using existing network resources. There is a limit of three clusters within each region, but there is no limit on the number of nodes and pods you can create within each cluster. This statement is incorrect. There is no specific limit on the number of clusters you can create within a region in OKE. However, there may be certain limits or quotas on resources that can impact the number of clusters you can create.

Question 2

Question Type: MultipleChoice

(CHK_1>3) You have an e-commerce application that loads customers' transactional data into the Oracle Cloud Infrastructure (OCI) Streaming service. The data must now be extracted and transformed before sending it to a third-party REST endpoint. You have been directed to leverage the OCI Service Connector Hub to automate this process. Which configuration option would address this requirement?

Options:

- A- Configure a new service connector as follows: * Source: Streaming * Task: Functions * Target: Functions
- B- Configure a new service connector as follows: * Source: Streaming * Task: API Gateway * Target: Notifications
- C- Configure a new service connector as follows: * Source: Streaming * Task: None * Target: Notifications
- D- Configure a new service connector as follows: * Source: Streaming * Task: API Gateway * Target: Functions
- E- Configure a new service connector as follows: * Source: Streaming * Task: Functions * Target: API Gateway

Answer:

C

Explanation:

To address the requirement of extracting and transforming data from the Oracle Cloud Infrastructure (OCI) Streaming service and sending it to a third-party REST endpoint using the OCI Service Connector Hub, the best configuration option is: Configure a new service connector as follows: * Source: Streaming * Task: None * Target: Notifications By selecting the Streaming service as the source, you can capture the transactional data from the stream. Since there is a need to transform and send the data to a third-party REST endpoint, you don't need to specify any specific task in the connector. The target is set to Notifications, which allows you to send the transformed data to an endpoint outside of the OCI environment. Notifications can be configured to deliver the data to various supported destinations, including HTTP endpoints, email addresses, and more. This configuration enables you to automate the process of extracting data from the streaming service and sending it to the desired third-party REST endpoint, fulfilling the requirement of extracting, transforming, and forwarding the data.

Question 3

Question Type: MultipleChoice

Which of the following TWO statements are TRUE about deleting a Kubernetes cluster? (Choose two.)

Options:

- A-** Upon deleting a cluster, other resources created during the cluster creation process or associated with the cluster (such as VCNS, Internet Gateways, NAT Gateways, Route Tables, Security Lists, B. Load Balancers, and Block Volumes) are deleted automatically.
- B-** If you change the auto-generated name of a worker node and then delete the cluster, the renamed worker node is not deleted.
- C-** Upon deleting a cluster, no other resources created during the cluster creation process or associated with the cluster (such as VCNS, Internet Gateways, NAT Gateways, Route Tables, Security Lists, Load Balancers, and Block Volumes) are deleted automatically.
- D-** Changing the auto-generated name of a worker node does not affect the deletion of the worker node when the cluster in which it is created is deleted.
- E-** You cannot change the autogenerated names of the worker nodes in the format oke-c---<slot> within a Kubernetes cluster.

Answer:

B, C

Explanation:

The correct statements about deleting a Kubernetes cluster are: If you change the auto-generated name of a worker node and then delete the cluster, the renamed worker node is not deleted. Changing the name of a worker node does not affect its deletion when the cluster is deleted. The cluster deletion process does not consider the renamed worker nodes and will delete all worker nodes associated with the cluster. Upon deleting a cluster, no other resources created during the cluster creation process or associated with the cluster (such as VCNs, Internet Gateways, NAT Gateways, Route Tables, Security Lists, Load Balancers, and Block Volumes) are deleted automatically. These additional resources are not automatically deleted when the cluster is deleted. You need to manage the deletion of these resources separately, if desired. Therefore, the correct statements are that the renamed worker nodes are not deleted when the cluster is deleted, and other associated resources are not automatically deleted when the cluster is deleted.

Question 4

Question Type: MultipleChoice

You encounter an unexpected error when invoking Oracle Functions from your Cloud Shell session named myfunction in the myapp application. Which option will get you more information on the error?

Options:

- A- `fn --verbose invoke myapp myfunction`
- B- `fn --debug invoke myapp myfunction`
- C- Contact Oracle support with your error message
- D- `DEBUG=1 fn invoke myapp myfunction`

Answer:

D

Explanation:

The option that will get you more information on the error when invoking Oracle Functions from your Cloud Shell session is: `'DEBUG=1 fn invoke myapp myfunction'`. Setting the environment variable `DEBUG=1` before invoking the function using the `fn` command allows you to enable debug mode, which provides more detailed information about the execution of the function. This can be useful for troubleshooting and understanding the root cause of the error. By using the command `'DEBUG=1 fn invoke myapp myfunction'`, the function invocation will be executed with debug mode enabled, and additional debug information will be displayed in the console output. This information can include stack traces, detailed error messages, and other relevant details that can help identify and resolve the issue. Using the verbose option (`--verbose`) or debug option (`--debug`) with the `fn` command may also provide additional information, but the specific behavior may depend on the version and configuration of the `fn` CLI tool. While contacting Oracle support with the error message is always an option, enabling debug mode using the `DEBUG=1` environment variable provides immediate access to more detailed information and can help in diagnosing and resolving the error more efficiently.

Question 5

Question Type: MultipleChoice

Which one of the following is NOT a valid backend-type supported by Oracle Cloud Infrastructure (OCI) API Gateway?

Options:

- A- STOCK_RESPONSE_BACKEND
- B- HTTP_BACKEND
- C- ORACLE_FUNCTIONS_BACKEND
- D- ORACLE_STREAMS_BACKEND

Answer:

D

Explanation:

Oracle Cloud Infrastructure (OCI) API Gateway supports various backend-types to handle incoming requests and route them to the appropriate backend service. However, 'ORACLE_STREAMS_BACKEND' is not a valid backend-type in OCI API Gateway.

'STOCK_RESPONSE_BACKEND' is a valid backend-type that allows you to configure static responses directly in the API Gateway

without routing requests to any specific backend service. This can be useful for handling simple requests or returning predefined responses. 'HTTP_BACKEND' is another valid backend-type that enables routing requests to an HTTP backend service. 'ORACLE_FUNCTIONS_BACKEND' is a valid backend-type that allows you to route requests to Oracle Functions, which are serverless functions that can be used to execute specific logic or operations. However, 'ORACLE_STREAMS_BACKEND' is not a valid backend-type in OCI API Gateway. The API Gateway does not have native support for Oracle Streams as a backend service. Oracle Streams is a feature of Oracle Database that provides a publish-subscribe mechanism for data replication and distribution within an Oracle Database environment. It is not directly integrated as a backend service in OCI API Gateway.

Question 6

Question Type: MultipleChoice

You are developing a serverless application with Oracle Functions and Oracle Cloud Infrastructure Object Storage. Your function needs to read a JSON file object from an Object Storage bucket named "input-bucket" in compartment "qa-compartment". Your corporate security standards mandate the use of Resource Principals for this use case. Which two statements are needed to implement this use case? (Choose two.)

Options:

A- Set up a policy to grant all functions read access to the bucket: allow all functions in compartment qa-compartment to read objects in target.bucket.name= 'input-bucket'

- B-** Set up a policy to grant your user account read access to the bucket: allow user XYZ to read objects in compartment qa-compartment where target.bucket.name= 'input-bucket'
- C-** Set up the following dynamic group for your function's OCID: Name: read-file-dg Rule: resource.id = 'ocid1.fnfunc.oc1.phx.aaaaaaaakeaobctakezjz5i4ujj7g25q7sx5mvr55pms6f4da'
- D-** No policies are needed. By default, every function has read access to Object Storage buckets in the tenancy.
- E-** Set up a policy with the following statement to grant read access to the bucket: allow dynamic-group read-file-dg to read objects in compartment qa- compartment where target.bucket.name= 'input-bucket'

Answer:

C, E

Explanation:

The correct answers are: Set up the following dynamic group for your function's OCID: Name: read-file-dg Rule: resource.id = 'ocid1.fnfunc.oc1.phx.aaaaaaaakeaobctakezjz5i4ujj7g25q7sx5mvr55pms6f4da' Set up a policy with the following statement to grant read access to the bucket: Statement: allow dynamic-group read-file-dg to read objects in compartment qa-compartment where target.bucket.name = 'input-bucket' To implement the use case of reading a JSON file object from an Object Storage bucket using Resource Principals with Oracle Functions, you need to configure the following: Create a dynamic group named 'read-file-dg' and associate it with your function's OCID. This dynamic group helps identify the function as a member of the group for policy enforcement. Create a policy that grants read access to the bucket. The policy statement should allow the dynamic group 'read-file-dg' to read objects in the compartment 'qa-compartment' and specify the target bucket name as 'input-bucket'. This policy ensures that the function has the necessary permissions to access the specified bucket. By setting up the dynamic group and policy, you ensure that the function, as a member of the dynamic group, has the required read access to the specified Object Storage bucket in the specified compartment.

To Get Premium Files for 1Z0-1084-23 Visit

<https://www.p2pexams.com/products/1z0-1084-23>

For More Free Questions Visit

<https://www.p2pexams.com/oracle/pdf/1z0-1084-23>

