



1Z0-1084-21^{Q&As}

Oracle Cloud Infrastructure Developer 2021 Associate

Pass Oracle 1Z0-1084-21 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass4itsure.com/1z0-1084-21.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Oracle
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





QUESTION 1

You are developing a serverless application with Oracle Functions. You have created a function in compartment named prod. When you try to invoke your function you get the following error. Error invoking function. status: 502 message: dhcp options ocid1.dhcpoptions.oc1.phx.aaaaaaac... does not exist or Oracle Functions is not authorized to use it How can you resolve this error?

- A. Create a policy:Allow function-family to use virtual-network-family in compartment prod
- B. Create a policy:Allow any-user to manage function-family and virtual-network-family in compartment prod
- C. Create a policy:Allow service FaaS to use virtual-network-family in compartment prod
- D. Deleting the function and redeploying it will fix the problem

Correct Answer: C

Invoking a function returns a FunctionInvokeSubnetNotAvailable message and a 502 error (due to a DHCP Options issue)

When you invoke a function that you've deployed to Oracle Functions, you might see the following error message:

```
{"code":"FunctionInvokeSubnetNotAvailable","message":"dhcp options ocid1.dhcpoptions..... does not exist or Oracle Functions is not authorized to use it"}
```

Fn: Error invoking function. status: 502 message: dhcp options ocid1.dhcpoptions..... does not exist or Oracle Functions is not authorized to use it If you see this error:

Double-check that a policy has been created to give Oracle Functions access to network resources.

Service Access to Network Resources

When Oracle Functions users create a function or application, they have to specify a VCN and a subnet in which to create them. To enable the Oracle Functions service to create the function or application in the specified VCN and subnet, you must create an identity policy to grant the Oracle Functions service access to the compartment to which the network resources belong. To create a policy to give the Oracle Functions service access to network resources:

Log in to the Console as a tenancy administrator.

Create a new policy in the root compartment:

Open the navigation menu. Under Governance and Administration, go to Identity and click Policies. Follow the instructions in To create a policy, and give the policy a name (for example, functions- service-networkaccess).



Specify a policy statement to give the Oracle Functions service access to the network resources in the compartment:

Allow service FaaS to use virtual-network-family in compartment For example:

Allow service FaaS to use virtual-network-family in compartment acme-network Click Create.

Double-check that the set of DHCP Options in the VCN specified for the application still exists.

QUESTION 2

Your Oracle Cloud Infrastructure Container Engine for Kubernetes (OKE) administrator has created an OKE cluster with one node pool in a public subnet. You have been asked to provide a log file from one of the nodes for troubleshooting purpose.

Which step should you take to obtain the log file?

- A. ssh into the node using public key.
- B. ssh into the nodes using private key.
- C. It is impossible since OKE is a managed Kubernetes service.
- D. Use the username open and password to login.

Correct Answer: B

Kubernetes cluster is a group of nodes. The nodes are the machines running applications. Each node can be a physical machine or a virtual machine. The node's capacity (its number of CPUs and amount of memory) is defined when the node is created. A cluster comprises: - one or more master nodes (for high availability, typically there will be a number of master nodes) - one or more worker nodes (sometimes known as minions) Connecting to Worker Nodes Using SSH If you provided a public SSH key when creating the node pool in a cluster, the public key is installed on all worker nodes in the cluster. On UNIX and UNIX-like platforms (including Solaris and Linux), you can then connect through SSH to the worker nodes using the ssh utility (an SSH client) to perform administrative tasks. Note the following instructions assume the UNIX machine you use to connect to the worker node: Has the ssh utility installed. Has access to the SSH private key file paired with the SSH public key that was specified when the cluster was created. How to connect to worker nodes using SSH depends on whether you specified public or private subnets for the worker nodes when defining the node pools in the cluster. Connecting to Worker Nodes in Public Subnets Using SSH Before you can connect to a worker node in a public subnet using SSH, you must define an ingress rule in the subnet's security list to allow SSH access. The ingress rule must allow access to port 22 on worker nodes from source 0.0.0.0/0 and any source port To connect to a worker node in a public subnet through SSH from a UNIX machine using the ssh utility: 1- Find out the IP address of the worker node to which you want to connect. You can do this in a number of ways: Using kubectl. If you haven't already done so, follow the steps to set up the cluster's kubeconfig configuration file and (if necessary) set the KUBECONFIG environment variable to point to the file. Note that you must set up your own kubeconfig file. You cannot access a cluster using a kubeconfig file that a different user set up. See Setting Up Cluster Access. Then in a terminal window, enter kubectl get nodes to see the public IP addresses of worker nodes in node pools in the cluster. Using the Console. In the Console, display the Cluster List page and then select the cluster to which the worker node belongs. On the Node Pools tab, click the name of the node pool to which the worker node belongs. On the Nodes tab, you see the public IP address of every worker node in the node pool. Using the REST API. Use the ListNodePools operation to see the public IP addresses of worker nodes in a node pool. 2- In the terminal window, enter ssh opc@ to connect to the worker node, where is the IP address of the worker node that you made a note of earlier. For example, you might enter ssh opc@192.0.2.254. Note that if the SSH private key is not stored in the file or in the path that the ssh



utility expects (for example, the ssh utility might expect the private key to be stored in `~/.ssh/id_rsa`), you must explicitly specify the private key filename and location in one of two ways: Use the `-i` option to specify the filename and location of the private key. For example, `ssh -i ~/.ssh/my_keys/my_host_key_filename opc@192.0.2.254` Add the private key filename and location to an SSH

configuration file, either the client configuration file (`~/.ssh/config`) if it exists, or the system-wide client

configuration file (`/etc/ssh/ssh_config`). For example, you might add the following:

```
Host 192.0.2.254 IdentityFile ~/.ssh/my_keys/my_host_key_filename
```

For more about the ssh utility's configuration file, enter `man ssh_config` Note also that permissions on the

private key file must allow you read/write/execute access, but prevent other users from accessing the file.

For example, to set appropriate permissions, you might enter `chmod 600 ~/.ssh/my_keys/`

`my_host_key_filename`. If permissions are not set correctly and the private key file is accessible to other

users, the ssh utility will simply ignore the private key file.

QUESTION 3

What are two of the main reasons you would choose to implement a serverless architecture?

- A. No need for integration testing
- B. Reduced operational cost
- C. Improved In-function state management
- D. Automatic horizontal scaling
- E. Easier to run long-running operations

Correct Answer: BD

QUESTION 4

In the sample Kubernetes manifest file below, what annotations should you add to create a private load balancer In oracle Cloud infrastructure Container Engine for Kubernetes?



```
apiVersion: v1
kind: Service
metadata:
  name: my-nginx-svc
  labels:
    app: nginx
  annotations:
    <Fill in>
spec:
  type: LoadBalancer
  ports:
    - port: 80
  selector:
    app: nginx
```

```
apiVersion: v1
kind: Service
metadata:
  name: my-nginx-svc
  labels:
    app: nginx
  annotations:
    <Fill in>
spec:
  type: LoadBalancer
  ports:
    - port: 80
  selector:
    app: nginx
```

A. service.beta.kubernetes.io/oci-load-balancer-private:"true"

B. service.beta.kubernetes.io/oci-load-balancer-private: "true" service.beta.kubernetes.io/oci-load-balancer-subnet1: "ocid1.subnet.oc1..aaaaa....vdfw"

C. service.beta.kubernetes.io/oci-load-balancer-internal: "true"

D. service.beta.kubernetes.io/oci-load-balancer-internal: "true" service.beta.kubernetes.io/oci-load-balancer-subnet1: "ocid1.subnet.oc1..aaaaa....vdfw"

Correct Answer: D

https://docs.cloud.oracle.com/en-us/iaas/Content/ContEng/Tasks/contengcreatingloadbalancer.htm?TocPath=Services%7CExample%20Network%20Resource%20Configuration%7CUpgrading%20the%20Version%20of%20Kubernetes%20Running%20on%20a%20Master%20Node%7C_____2 Creating Internal Load Balancers in Public and Private Subnets You can create Oracle Cloud Infrastructure load balancers to control access to



services running on a cluster: When you create a `custom` cluster, you select an existing VCN that contains the network resources to be used by the new cluster. If you want to use load balancers to control traffic into the VCN, you select existing public or private subnets in that VCN to host the load balancers. When you create a `quick cluster`, the VCN that's automatically created contains a public regional subnet to host a load balancer. If you want to host load balancers in private subnets, you can add private subnets to the VCN later.

Alternatively, you can create an internal load balancer service in a cluster to enable other programs running in the same VCN as the cluster to access services in the cluster. You can host internal load balancers in public subnets and private subnets. To create an internal load balancer hosted on a public subnet, add the following annotation in the metadata section of the manifest file: `service.beta.kubernetes.io/oci-load-balancer-internal: "true"` To create an internal load balancer hosted on a private subnet, add both following annotations in the metadata section of the manifest file: `service.beta.kubernetes.io/oci-load-balancer-internal: "true"` `service.beta.kubernetes.io/oci-load-balancersubnet1: "ocid1.subnet.oc1..aaaaa....vdfw"` where `ocid1.subnet.oc1..aaaaa....vdfw` is the OCID of the private subnet.

QUESTION 5

A developer using Oracle Cloud Infrastructure (OCI) API Gateway must authenticate the API requests to their web application. The authentication process must be implemented using a custom scheme which accepts string parameters from the API caller. Which method can the developer use In this scenario?

- A. Create an authorizer function using request header authorization.
- B. Create an authorizer function using token-based authorization.
- C. Create a cross account functions authorizer.
- D. Create an authorizer function using OCI Identity and Access Management based authentication

Correct Answer: B

Having deployed the authorizer function, you enable authentication and authorization for an API deployment by including two different kinds of request policy in the API deployment specification: An authentication request policy for the entire API deployment that specifies: The OCID of the authorizer function that you deployed to Oracle Functions that will perform authentication and authorization. The request attributes to pass to the authorizer function. Whether unauthenticated callers can access routes in the API deployment. An authorization request policy for each route that specifies the operations a caller is allowed to perform, based on the caller's access scopes as returned by the authorizer function. Using the Console to Add Authentication and Authorization Request Policies To add authentication and authorization request policies to an API deployment specification using the Console: Create or update an API deployment using the Console, select the From Scratch option, and enter details on the Basic Information page. For more information, see Deploying an API on an API Gateway by Creating an API Deployment and Updating API Gateways and API Deployments. In the API Request Policies section of the Basic Information page, click the Add button beside Authentication and specify: Application in : The name of the application in Oracle Functions that contains the authorizer function. You can select an application from a different compartment. Function Name: The name of the authorizer function in Oracle Functions. Authentication Token: Whether the access token is contained in a request header or a query parameter. Authentication Token Value: Depending on whether the access token is contained in a request header or a query parameter, specify:

Header Name: If the access token is contained in a request header, enter the name of the header. Parameter Name: If the access token is contained in a query parameter, enter the name of the query parameter.

<https://docs.cloud.oracle.com/en-us/iaas/Content/APIGateway/Tasks/apigatewayaddingauthzauthn.htm>