

[Q16-Q40 Latest 1Z0-1084-21 Exam with Accurate Oracle Cloud Infrastructure Developer 2021 Associate PDF Questions [Dec 13, 2022]

[Dec 13, 2022] Latest 1Z0-1084-21 Exam with Accurate Oracle Cloud Infrastructure Developer 2021 Associate PDF Questions
Practice To 1Z0-1084-21 - DumpsMaterials Remarkable Practice On your Oracle Cloud Infrastructure Developer 2021 Associate Exam

Oracle 1Z0-1084-21 Exam Syllabus Topics:

TopicDetailsTopic 1- Explain the Challenges of testing in Cloud Native world- Describe Service Communication PatternsTopic 2- Overcome Security challenges with Cloud Native- Discuss the Meaning of Serverless ComputingTopic 3- Securing Cloud Native Applications- Use the Defense-in-depth approachTopic 4- Perform Tasks around Monitoring, Observability, and Alerting- Explain Distributed ComputingTopic 5- Develop Serverless Application with Oracle Functions- Develop Microservices and Applications for OKE

QUESTION 16

In a Linux environment, what is the default locations of the configuration file that Oracle Cloud Infrastructure CLI uses for profile information/

- * /etc/.oci/config
- * /usr/local/bin/config
- * \$HOME/.oci/config
- * /usr/bin/oci/config

Explanation

By default, the Oracle Cloud Infrastructure CLI configuration file is located at ~/.oci/config.

You might already have a configuration file as a result of installing the Oracle Cloud Infrastructure CLI.

QUESTION 17

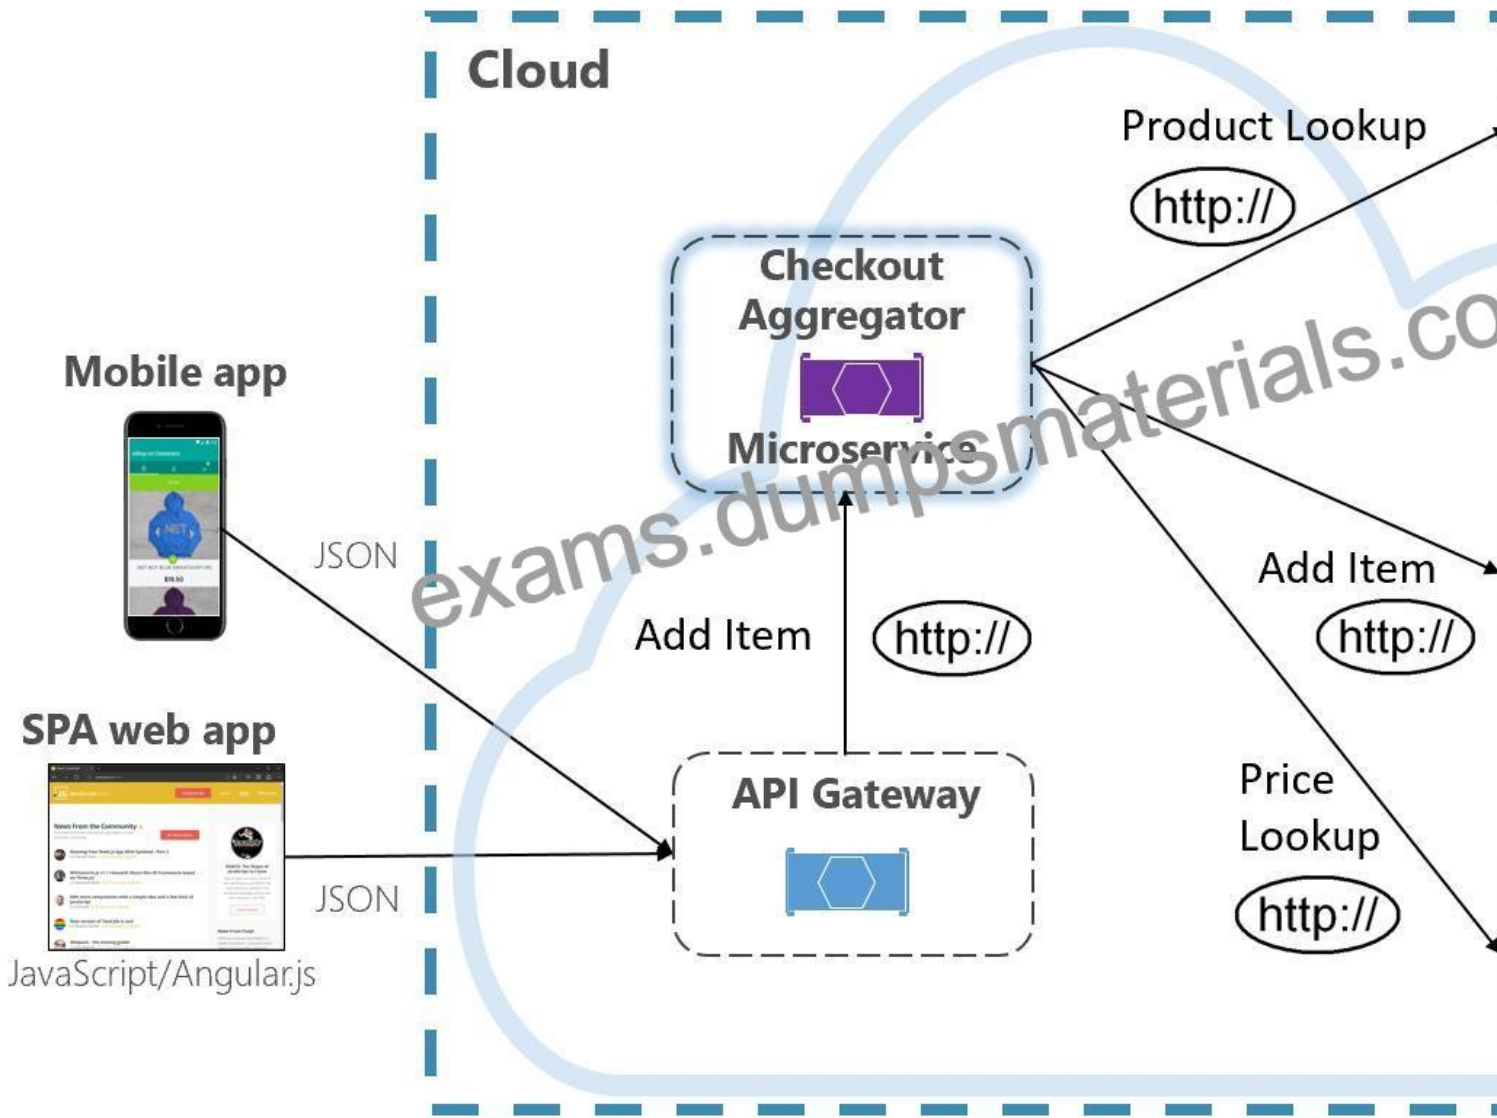
Which one of the statements describes a service aggregator pattern?

- * It is implemented in each service separately and uses a streaming service
- * It involves implementing a separate service that makes multiple calls to other backend services
- * It uses a queue on both sides of the service communication
- * It involves sending events through a message broker

Explanation

Service Aggregator Pattern

Another option for eliminating microservice-to-microservice coupling is an Aggregator microservice, shown in purple in Figure 4-10.



The pattern isolates an operation that makes calls to multiple back-end microservices, centralizing its logic into a specialized microservice. The purple checkout aggregator microservice in the previous figure orchestrates the workflow for the Checkout operation. It includes calls to several back-end microservices in a sequenced order. Data from the workflow is aggregated and returned to the caller. While it still implements direct HTTP calls, the aggregator microservice reduces direct dependencies among back-end microservices.

References:

<https://docs.microsoft.com/en-us/dotnet/architecture/cloud-native/service-to-service-communication#:~:text=Ser>

QUESTION 18

You need to execute a script on a remote instance through Oracle Cloud Infrastructure Resource Manager.

Which option can you use?

* Use `/bin/sh` with the full path to the location of the script to execute the script.

- * It cannot be done.
- * Download the script to a local desktop and execute the script.
- * Use remote-exec

Explanation

Using Remote Exec

With Resource Manager, you can use Terraform's remote exec functionality to execute scripts or commands on a remote computer. You can also use this technique for other provisioners that require access to the remote resource.

References:

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

QUESTION 19

As a cloud-native developer, you have written a web service for your company. You have used Oracle Cloud Infrastructure (OCI) API Gateway service to expose the HTTP backend. However, your security team has suggested that your web service should handle Distributed Denial-of-Service (DDoS) attack. You are time-constrained and you need to make sure that this is implemented as soon as possible.

What should you do in this scenario?

- * Use OCI virtual cloud network (VCN) segregation to control DDoS.
- * Use a third party service integration to implement a DDoS attack mitigation,
- * Use OCI API Gateway service and configure rate limiting.
- * Re-write your web service and implement rate limiting.

Explanation

Having created an API gateway and deployed one or more APIs on it, you'll typically want to limit the rate at which front-end clients can make requests to back-end services. For example, to:

• maintain high availability and fair use of resources by protecting back ends from being overwhelmed by too many requests

• prevent denial-of-service attacks

• constrain costs of resource consumption

• restrict usage of APIs by your customers/users in order to monetize APIs You apply a rate limit globally to all routes in an API deployment specification.

If a request is denied because the rate limit has been exceeded, the response header specifies when the request can be retried.

You can add a rate-limiting request policy to an API deployment specification by:

using the Console

editing a JSON file

```
{
```

```
&#8220;requestPolicies&#8221;: {  
  
&#8220;rateLimiting&#8221;: {  
  
&#8220;rateKey&#8221;: &#8220;CLIENT_IP&#8221;,  
  
&#8220;rateInRequestsPerSecond&#8221;: 10  
  
}  
  
},  
  
&#8220;routes&#8221;:[  
  
{  
  
&#8220;path&#8221;: &#8220;/hello&#8221;,  
  
&#8220;methods&#8221;: [&#8220;GET&#8221;],  
  
&#8220;backend&#8221;: {  
  
&#8220;type&#8221;: &#8220;ORACLE_FUNCTIONS_BACKEND&#8221;,  
  
&#8220;functionId&#8221;: &#8220;ocid1.fnfunc.oc1.phx.aaaaaaaab_____xmq&#8221;,  
  
}  
  
}  
  
]  
  
}
```

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

QUESTION 20

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

- * STOCK_RESPONSE_BACKEND
- * ORACLE_FUNCTIONS_BACKEND
- * ORACLE_STREAMS_BACKEND
- * HTTP_BACKEND

Explanation

In the API Gateway service, a back end is the means by which a gateway routes requests to the back-end services that implement APIs. If you add a private endpoint back end to an API gateway, you give the API gateway access to the VCN associated with that private endpoint.

You can also grant an API gateway access to other Oracle Cloud Infrastructure services as back ends. For example, you could grant

an API gateway access to Oracle Functions, so you can create and deploy an API that is backed by a serverless function.

API Gateway service to create an API gateway, you can create an API deployment to access HTTP and HTTPS URLs.

<https://docs.cloud.oracle.com/en-us/iaas/Content/APIGateway/Tasks/apigatewayusinghttpbackend.htm> API Gateway service to create an API gateway, you can create an API deployment that invokes serverless functions defined in Oracle Functions.

<https://docs.cloud.oracle.com/en-us/iaas/Content/APIGateway/Tasks/apigatewayusingfunctionsbackend.htm> API Gateway service, you can define a path to a stock response back end

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

QUESTION 21

In order to effectively test your cloud-native applications, you might utilize separate environments (development, testing, staging, production, etc.). Which Oracle Cloud Infrastructure (OCI) service can you use to create and manage your infrastructure?

- * OCI Compute
- * OCI Container Engine for Kubernetes
- * OCI Resource Manager
- * OCI API Gateway

Explanation

Resource Manager is an Oracle Cloud Infrastructure service that allows you to automate the process of provisioning your Oracle Cloud Infrastructure resources. Using Terraform, Resource Manager helps you install, configure, and manage resources through the `infrastructure-as-code` model.

References:

<https://docs.cloud.oracle.com/iaas/Content/ResourceManager/Concepts/resourcemanager.htm>

QUESTION 22

You are developing a polyglot serverless application using Oracle Functions. Which language cannot be used to write your function code?

- * PL/SQL
- * Python
- * Node.js
- * Java

Explanation

Overview of Functions:

The serverless and elastic architecture of Oracle Functions means there's no infrastructure administration or software administration for you to perform. You don't provision or maintain compute instances, and operating system software patches and upgrades are applied automatically. Oracle Functions simply ensures your app is highly-available, scalable, secure, and monitored. With Oracle Functions, you can write code in Java, Python, Node, Go, and Ruby (and for advanced use cases, bring your own Dockerfile, and Graal VM). You can then deploy your code, call it directly or trigger it in response to events, and get billed only for the resources consumed during the execution.

References:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Functions/Concepts/functionoverview.htm>

QUESTION 23

Which two statements accurately describe Oracle SQL Developer Web on Oracle Cloud Infrastructure (OCI) Autonomous Database?

- * It is available for databases with dedicated Exadata infrastructure only.
- * After provisioning into an OCI compute Instance, it can automatically connect to the OCI Autonomous Databases instances.
- * It is available for databases with both dedicated and shared Exadata infrastructure.
- * It provides a development environment and a data modeler interface for OCI Autonomous Databases.
- * It must be enabled via OCI Identity and Access Management policy to get access to the Autonomous Databases instances.

Explanation

Oracle SQL Developer Web Oracle SQL Developer Web in Autonomous Data Warehouse provides a development environment and a data modeler interface for Autonomous Databases. SQL Developer Web is available for databases with both dedicated Exadata infrastructure and shared Exadata infrastructure.

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

QUESTION 24

You are developing a distributed application and you need a call to a path to always return a specific JSON content deploy an Oracle Cloud Infrastructure API Gateway with the below API deployment specification.

```
{
  "routes": [{
    "path": "/hello",
    "methods": ["GET"],
    "backend": {
      "type": "STOCK_RESPONSE_BACKEND",
      "status": 200,
      "headers": [{
        "name": "Content-Type",
        "value": "application/json"
      }],
      "body": "{\n\"myjson\": \"consistent response\"}"
    }
  ]
}
```

What is the correct value for type?

- * STOCK_RESPONSE_BACKEND
- * CONSTANT_BACKEND
- * JSON_BACKEND
- * HTTP_BACKEND

Explanation

Adding Stock Responses as an API Gateway Back End:

You often want to verify that an API has been successfully deployed on an API gateway without having to set up an actual back-end service. One approach is to define a route in the API deployment specification that has a path to a

‘dummy’ back end. On receiving a request to that path, the API gateway itself acts as the back end and returns a stock response you’ve specified.

Equally, there are some situations in a production deployment where you’ll want a particular path for a route to consistently return the same stock response without sending a request to a back end. For example, when you want a call to a path to always return a specific HTTP status code in the response.

Using the API Gateway service, you can define a path to a stock response backend that always returns the same:

HTTP status code

HTTP header fields (name-value pairs)

content in the body of the response

“type”: “STOCK_RESPONSE_BACKEND” indicates that the API gateway itself will act as the back end and return the stock response you define (the status code, the header fields and the body content).

References:

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

QUESTION 25

Which is NOT a supported SDK Oracle Cloud Infrastructure (OCI)?

- * Go SDK
- * Java SDK
- * NET SDK
- * Ruby SDK
- * Python SDK

Explanation

<https://docs.cloud.oracle.com/en-us/iaas/Content/API/Concepts/sdks.htm>

* Software Development Kits (SDKs) Build and deploy apps that integrate with Oracle Cloud Infrastructure services. Each SDK provides the tools you need to develop an app, including code samples and documentation to create, test, and troubleshoot. In addition, if you want to contribute to the development of the SDKs, they are all open source and available on GitHub.

* SDK for Java

* Python SDK

* Ruby SDK

* Go SDK

<https://docs.cloud.oracle.com/en-us/iaas/Content/API/Concepts/sdkconfig.htm>

QUESTION 26

Which two are characteristics of microservices?

- * Microservices are hard to test in isolation.
- * Microservices can be independently deployed.
- * All microservices share a data store.
- * Microservices can be implemented in limited number of programming languages.
- * Microservices communicate over lightweight APIs.

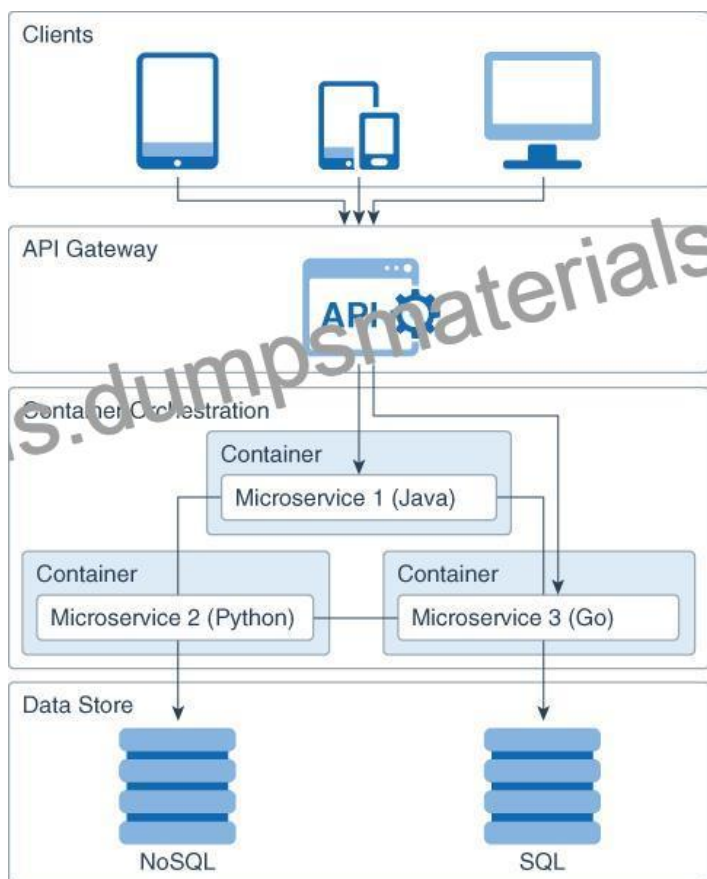
Explanation

Learn About the Microservices Architecture

If you want to design an application that is multilanguage, easily scalable, easy to maintain and deploy, highly available, and that minimizes failures, then use the microservices architecture to design and deploy a cloud application.

In a microservices architecture, each microservice owns a simple task, and communicates with the clients or with other microservices by using lightweight communication mechanisms such as REST API requests.

The following diagram shows the architecture of an application that consists of multiple microservices.



Microservices enable you to design your application as a collection of loosely coupled services. Microservices follow the share-nothing model, and run as stateless processes. This approach makes it easier to scale and maintain the application.

The API layer is the entry point for all the client requests to a microservice. The API layer also enables the microservices to communicate with each other over HTTP, gRPC, and TCP/UDP.

The logic layer focuses on a single business task, minimizing the dependencies on the other microservices.


This layer can be written in a different language for each microservice.

The data store layer provides a persistence mechanism, such as a database storage engine, log files, and so on.

Consider using a separate persistent data store for each microservice.

Typically, each microservice runs in a container that provides a lightweight runtime environment.

Loosely coupled with other services – enables a team to work independently the majority of time on their service(s) without being impacted by changes to other services and without affecting other services



References:

<https://docs.oracle.com/en/solutions/learn-architect-microservice/index.html>

<https://microservices.io/patterns/microservices.html>

<https://www.techjini.com/blog/microservices/>

QUESTION 27

What is one of the differences between a microservice and a serverless function?

- * Microservices are used for long running operations and serverless functions for short running operations.
- * Microservices always use a data store and serverless functions never use a data store.
- * Microservices are stateless and serverless functions are stateful.
- * Microservices are triggered by events and serverless functions are not.

Explanation

microservice is larger and can do more than a function. A function is a relatively small bit of code that performs only one action in response to an event.

Many microservices can run on several servers, and different instances of a specific microservice can run on several servers.

In many cases, microservices can be decomposed into a number of smaller stateless functions. The difference between microservices and functions is not simply the size. Functions are stateless, and they require no knowledge about or configuration of the underlying server-hence, the term serverless.

Microservices are best suited for long-running, complex applications that have significant resource and management requirements. You can migrate an existing monolithic application to microservices, which makes it easier to modularly develop features for the application and deploy it in the cloud. Microservices are also a good choice for building e-commerce sites, as they can retain information throughout a transaction and meet the needs of a 24/7 customer base.

On the other hand, serverless functions only execute when needed. Once the execution is over, the computing instance that runs the code decommissions itself. Serverless aligns with applications that are event driven, especially when the events are sporadic and the event processing is not resource-intensive. Serverless is a good choice when developers need to deploy fast and there are minimal application scaling concerns. For example, a good use of serverless computing is a scheduled task that needs to perform some data aggregation and will execute for just a few seconds.



References:

<https://www.cloudflare.com/learning/serverless/glossary/serverless-microservice/>

<https://developer.oracle.com/java/fn-project-introduction.html>

<https://searcharchitecture.techtarget.com/answer/When-should-I-choose-between-serverless-and-microservice>

QUESTION 28

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?

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

Explanation

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@<node_ip_address>` to connect to the worker node, where <node_ip_address> 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.

References:

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

QUESTION 29

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?

* 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 *
```

* Set up the following dynamic group for your function's OCID: Name: `read-file-dg` Rule: `resource.id = ; ocid1.fnunc.ocl-phx.aaaaaaakeaobctakezjz5i4ujj7g25q7sx5mvr55pms6f`

4da !

* 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;
```

* 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;
```

* No policies are needed. By default, every function has read access to Object Storage buckets in the tenancy

Explanation

When a function you've deployed to Oracle Functions is running, it can access other Oracle Cloud Infrastructure resources. For example:

– You might want a function to get a list of VCNs from the Networking service.

– You might want a function to read data from an Object Storage bucket, perform some operation on the data, and then write the modified data back to the Object Storage bucket.

To enable a function to access another Oracle Cloud Infrastructure resource, you have to include the function in a dynamic group, and then create a policy to grant the dynamic group access to that resource.

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

QUESTION 30

You created a pod called “nginx” and its state is set to Pending.

Which command can you run to see the reason why the “nginx” pod is in the pending state?

- * kubectl logs pod nginx
- * kubectl describe pod nginx
- * kubectl get pod nginx
- * Through the Oracle Cloud Infrastructure Console

Explanation

Debugging Pods

The first step in debugging a pod is taking a look at it. Check the current state of the pod and recent events with the following command:

```
kubectl describe pods ${POD_NAME}
```

Look at the state of the containers in the pod. Are they all Running? Have there been recent restarts?

Continue debugging depending on the state of the pods.

My pod stays pending

If a pod is stuck in Pending it means that it can not be scheduled onto a node. Generally this is because there are insufficient resources of one type or another that prevent scheduling. Look at the output of the `kubectl describe` … command above. There should be messages from the scheduler about why it can not schedule your pod.

<https://kubernetes.io/docs/tasks/debug-application-cluster/debug-pod-replication-controller/>

QUESTION 31

You are tasked with developing an application that requires the use of Oracle Cloud Infrastructure (OCI) APIs to POST messages to a stream in the OCI Streaming service.

Which statement is incorrect?

- * The request must include an authorization signing string including (but not limited to) `x-content-sha256`, `content-type`, and `content-length` headers.
- * The `Content-Type` header must be Set to `application/json`

- * An HTTP 401 will be returned if the client's clock is skewed more than 5 minutes from the server's.
- * The request does not require an Authorization header.

Explanation

Authorization Header

The Oracle Cloud Infrastructure signature uses the Signature Authentication scheme (with an Authorization header), and not the Signature HTTP header.

Required Credentials and OCIDs

You need an API signing key in the correct format. See Required Keys and OCIDs.

You also need the OCIDs for your tenancy and user. See Where to Get the Tenancy's OCID and User's OCID.

Summary of Signing Steps

In general, these are the steps required to sign a request:

Form the HTTPS request (SSL protocol/TLS 1.2 is required).

Create the signing string, which is based on parts of the request.

Create the signature from the signing string, using your private key and the RSA-SHA256 algorithm.

Add the resulting signature and other required information to the Authorization header in the request.

References:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Streaming/Concepts/streamingoverview.htm>

<https://docs.cloud.oracle.com/en-us/iaas/Content/API/Concepts/signingrequests.htm>

QUESTION 32

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?

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

Explanation

Using Authorizer Functions to Add Authentication and Authorization to API Deployments:

You can control access to APIs you deploy to API gateways using an authorizer function; (as described in this topic), or using JWTs (as described in Using JSON Web Tokens (JWTs) to Add Authentication and Authorization to API Deployments).

You can add authentication and authorization functionality to API gateways by writing an `authorizer function`; that:

1. Processes request attributes to verify the identity of a caller with an identity provider.
2. Determines the operations that the caller is allowed to perform.
3. Returns the operations the caller is allowed to perform as a list of `access scopes`; (an `access scope` is an arbitrary string used to determine access).


Optionally returns a key-value pair for use by the API deployment. For example, as a context variable for use in an HTTP back end definition (see Adding Context Variables to Policies and HTTP Back End Definitions).

Create an `authorizer function` using request header authorization implemented using a custom scheme which accepts string parameters from the API caller.

Managing Input Parameters

In our case we will need to manage quite a few static parameters in our code. For example the URLs of the secrets service endpoints, the username and other constant parameterised data. We can manage these either at Application or Function level (an OCI Function is packaged in an Application which can contain multiple Functions). In this case I will create function level parameters. You can use the following command to create the parameters:

```
fn config function test idcs-assert idcsClientId aedc15531bc8xxxxxxxxxxbd8a193
```



References:

<https://technology.amis.nl/2020/01/03/oracle-cloud-api-gateway-using-an-authorizer-function-for-client-secret-a>

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

<https://www.ateam-oracle.com/how-to-implement-an-oci-api-gateway-authorization-fn-in-nodejs-that-accesses-o>

QUESTION 33

What is the communication method between different Cloud native applications services?

- * Complex and asynchronous
- * Basic and synchronous
- * Complex and synchronous
- * Basic and asynchronous

Explanation

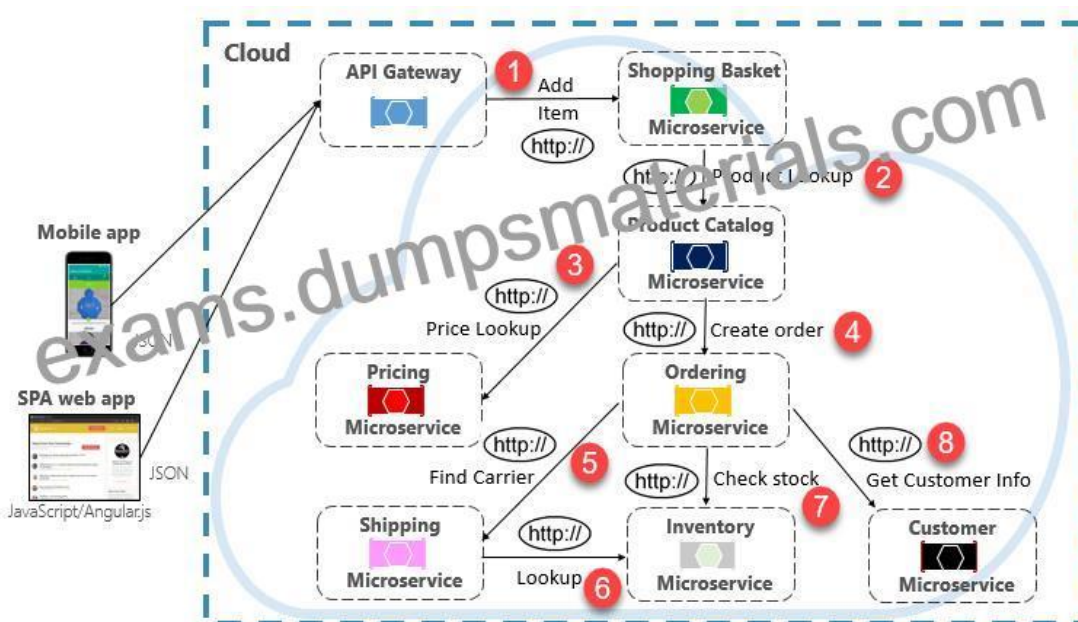
What Is Cloud Native?

Cloud native technologies are characterized by the use of containers, microservices, serverless functions, development pipelines, infrastructure expressed as code, event-driven applications, and Application Programming Interfaces (APIs). Cloud native enables faster software development and the ability to build applications that are resilient, manageable, observable, and dynamically scalable to global enterprise levels.

When constructing a cloud-native application, you'll want to be sensitive to how back-end services communicate with each other. Ideally, the less inter-service communication, the better. However, avoidance isn't always possible as back-end services often rely on one another to complete an operation.

While direct HTTP calls between microservices are relatively simple to implement, care should be taken to minimize this practice. To start, these calls are always synchronous and will block the operation until a result is returned or the request times out. What were once self-contained, independent services, able to evolve independently and deploy frequently, now become coupled to each other. As coupling among microservices increase, their architectural benefits diminish.

Executing an infrequent request that makes a single direct HTTP call to another microservice might be acceptable for some systems. However, high-volume calls that invoke direct HTTP calls to multiple microservices aren't advisable. They can increase latency and negatively impact the performance, scalability, and availability of your system. Even worse, a long series of direct HTTP communication can lead to deep and complex chains of synchronous microservices calls, shown in Figure 4-9:



A message queue is an intermediary construct through which a producer and consumer pass a message.

Queues implement an asynchronous, point-to-point messaging pattern.

Events

Message queuing is an effective way to implement communication where a producer can asynchronously send a consumer a message.

References:

<https://www.xenonstack.com/blog/cloud-native-architecture/>

<https://www.oracle.com/sa/cloud/cloud-native/>

<https://www.oracle.com/technetwork/topics/entarch/cloud-native-app-development-wp-3664668.pdf>

QUESTION 34

Which two handle Oracle Functions authentication automatically?

- * Oracle Cloud Infrastructure SDK
- * cURL
- * Oracle Cloud Infrastructure CLI
- * Signed HTTP Request
- * Fn Project CLI

Explanation

If you use the Fn Project CLI or the Oracle Cloud Infrastructure CLI to invoke a function, authentication is handled for you. See [Using the Fn Project CLI to Invoke Functions](#) and [Using the Oracle Cloud Infrastructure CLI to Invoke Functions](#).

If you use an Oracle Cloud Infrastructure SDK to invoke a function, you can use the SDK to handle authentication. See [Using SDKs to Invoke Functions](#).

If you make a signed HTTP request to a function's invoke endpoint, you'll have to handle authentication yourself by including a signature and the OCID of the compartment to which the function belongs in the request header. Fn Project CLI you can create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure and specify `provider oracle`. This option enables Oracle Functions to perform authentication and authorization using Oracle Cloud Infrastructure request signing, private keys, user groups, and policies that grant permissions to those user groups.

References:

<https://blogs.oracle.com/developers/oracle-functions-invoking-functions-automatically-with-cloud-events>

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

QUESTION 35

Which is NOT a supported SDK on Oracle Cloud Infrastructure (OCI)?

- * Ruby SDK
- * Java SDK
- * Python SDK
- * Go SDK
- * .NET SDK

Explanation

<https://docs.cloud.oracle.com/en-us/iaas/Content/API/Concepts/sdks.htm>

QUESTION 36

Which header is NOT required when signing GET requests to Oracle Cloud Infrastructure APIs?

- * date or x-date
- * (request-target)
- * content-type
- * host

Explanation

Authorization Header

The Oracle Cloud Infrastructure signature uses the `Signature`;Authentication scheme (with an Authorization header), and not the Signature HTTP header.

This section describes the headers that must be included in the signing string:

For GET and DELETE requests (when there's no content in the request body), the signingstring must include at least these headers:

(request-target)

host

date or x-date (if both are included, Oracle uses x-date)

For PUT and POST requests (when there's content in the request body), the signing string must include at least these headers:

(request-target)

host

date or x-date (if both are included, Oracle uses x-date)

x-content-sha256 (except for Object Storage PUT requests; see the next section) content-type content-length References:

<https://docs.cloud.oracle.com/en-us/iaas/Content/API/Concepts/signingrequests.htm>

QUESTION 37

How do you perform a rolling update in Kubernetes?

- * kubectl rolling-update
- * kubectl upgrade <deployment-name> -image=*image:v2
- * kubectl update -c <container>
- * kubectl rolling-update <deployment-name> -image=image

Explanation

<https://docs.oracle.com/en/cloud/iaas/wercker-cloud/wercm/quickstarts/platforms/kubernetes/>

QUESTION 38

Which two statements accurately describe an Oracle Functions application?

- * A small block of code invoked in response to an Oracle Cloud Infrastructure (OCI) Events service
- * A Docker image containing all the functions that share the same configuration
- * An application based on Oracle Functions, Oracle Cloud Infrastructure (OCI) Events and OCI API Gateway services
- * A common context to store configuration variables that are available to all functions in the application
- * A logical group of functions

Explanation

Oracle Functions Concepts:

This topic describes key concepts you need to understand when using Oracle Functions.

Applications:

In Oracle Functions, an application is:

1. a logical grouping of functions
2. a common context to store configuration variables that are available to all functions in the application
3. a way to ensure function runtime isolation

When you define an application in Oracle Functions, you specify the subnets in which to run the functions in the application. When functions from different applications are invoked simultaneously, Oracle Functions ensures these function executions are isolated from each other.

Oracle Functions shows applications and their functions in the Console.

References:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Functions/Concepts/functionsconcepts.htm>

QUESTION 39

You are using Oracle Cloud Infrastructure (OCI) Resource Manager to manage your infrastructure lifecycle and wish to receive an email each time a Terraform action begins.

How should you use the OCI Events service to do this without writing any code?

- * Create an OCI Notifications topic and email subscription with the destination email address. Then create an OCI Events rule matching Resource Manager Stack Update; condition, and select the notification topic for the

corresponding action.

* Create an OCI Notification topic and email subscription with the destination email address. Then create an OCI Events rule matching Resource Manager job Create; condition, and select the notification topic for the corresponding action.

* Create a rule in OCI Events service matching the Resource Manager Stack Update; condition. Then select Action Type: Email; and provide the destination email address.

* Create an OCI Email Delivery configuration with the destination email address. Then create an OCI Events rule matching Resource Manager Job Create; condition, and select the email configuration for the corresponding action.

Managing Topics and Subscriptions:


1. Create Notifications Topic and Subscription

If a suitable Notifications topic doesn't already exist, then you must log in to the Console as a tenancy administrator and create it. Whether you use an existing topic or create a new one, add an email address as a subscription so that you can monitor that email account for notifications

2. Using the Console to Create a Rule

Use the Console to create a rule with a pattern that matches bucket creation events emitted by Object Storage.

Specify the Notifications topic you created as an action to deliver matching events. To test your rule, create a bucket. Object Storage emits an event which triggers the action. Check the email specified in the subscription to receive your notification



References:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Events/Concepts/eventsgetstarted.htm>

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

<https://docs.cloud.oracle.com/en-us/iaas/Content/Notification/Concepts/notificationoverview.htm>

<https://docs.cloud.oracle.com/en-us/iaas/Content/Events/Concepts/filterevents.htm>

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

QUESTION 40

Which two are required to enable Oracle Cloud Infrastructure (OCI) Container Engine for Kubernetes (OKE) cluster access from the kubectl CLI?

- * An SSH key pair with the public key added to cluster worker nodes
- * Install and configure the OCI CLI
- * OCI Identity and Access Management Auth Token
- * Tiller enabled on the OKE cluster
- * A configured OCI API signing key pair

Explanation

Setting Up Local Access to Clusters

To set up a kubeconfig file to enable access to a cluster using a local installation of kubectl and the Kubernetes Dashboard:

Step 1: Generate an API signing key pair

Step 2: Upload the public key of the API signing key pair

Step 3: Install and configure the Oracle Cloud Infrastructure CLI

Step 4: Set up the kubeconfig file

Step 5: Verify that kubectl can access the cluster

References:

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

Exam Questions and Answers for 1Z0-1084-21 Study Guide Questions and Answers!:

<https://www.dumpsmaterials.com/1Z0-1084-21-real-torrent.html>