

Oracle

1Z0-1084-21 Exam

Oracle Cloud Infrastructure Developer 2021 Associate

**Questions & Answers
Demo**

Version: 8.0

Question: 1

Which one of the statements describes a service aggregator pattern?

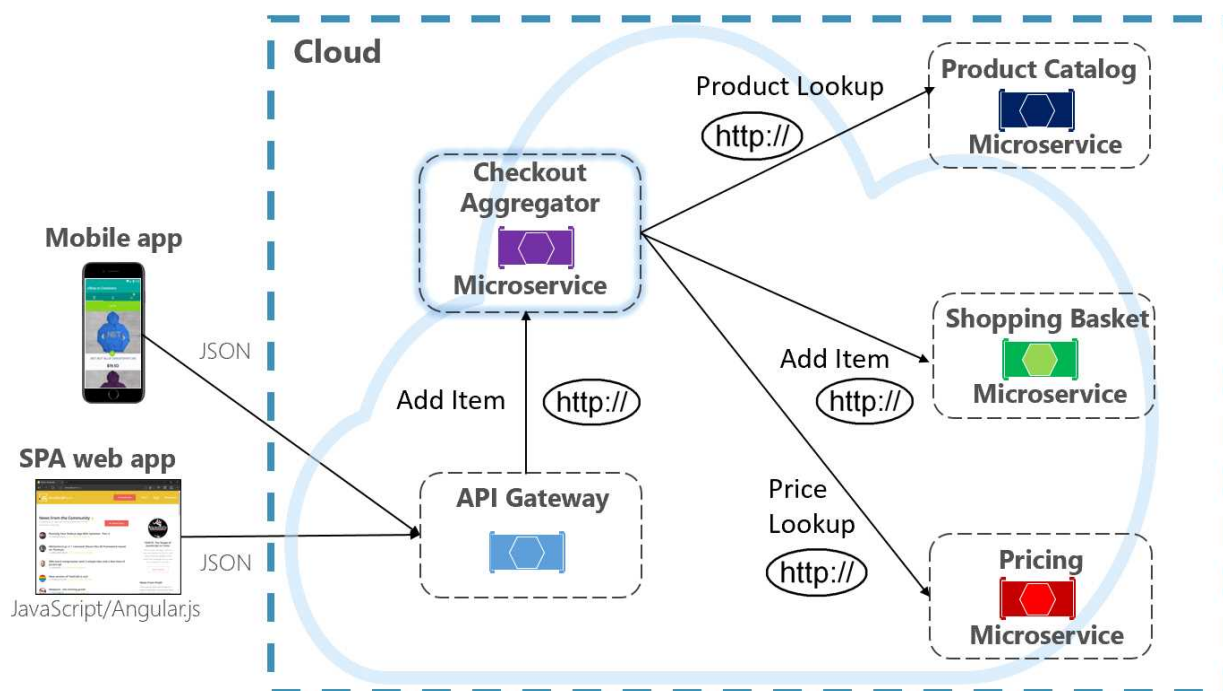
- A. It is implemented in each service separately and uses a streaming service
- B. It involves implementing a separate service that makes multiple calls to other backend services
- C. It uses a queue on both sides of the service communication
- D. It involves sending events through a message broker

Answer: B

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.

Reference:

<https://docs.microsoft.com/en-us/dotnet/architecture/cloud-native/service-to-service-communication#:~:text=Service%20Aggregator%20Pattern&text=The%20pattern%20isolates%20an%20operation,logic%20into%20a%20specialized%20microservice.&text=While%20it%20still%20implements%20direct,dependencies%20among%20back%2Dend%20microservices.>

Question: 2

Which two statements accurately describe an Oracle Functions application?

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

Answer: DE

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.

Reference:

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

Question: 3

You have been asked to create a stateful application deployed in Oracle Cloud Infrastructure (OCI) Container Engine for Kubernetes (OKE) that requires all of your worker nodes to mount and write data to persistent volumes.

Which two OCI storage services should you use?

- A. Use OCI File Services as persistent volume.
- B. Use GlusterFS as persistent volume.
- C. Use OCI Block Volume backed persistent volume.
- D. Use open source storage solutions on top of OCI.

E. Use OCI Object Storage as persistent volume.

Answer: AC

Explanation:


A PersistentVolume (PV) is a piece of storage in the cluster that has been provisioned by an administrator. PVs are volume plugins like Volumes, but have a lifecycle independent of any individual Pod that uses the PV.

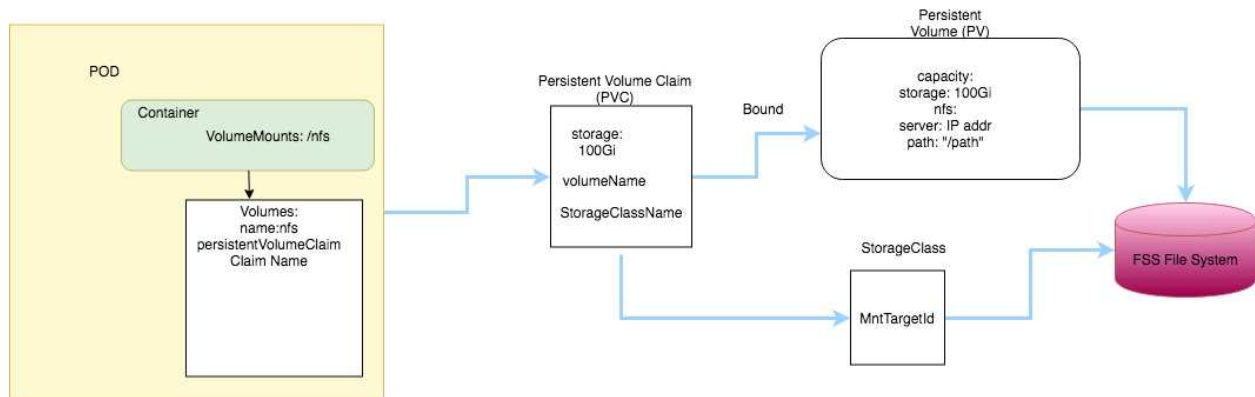
A PersistentVolumeClaim (PVC) is a request for storage by a user. It is similar to a Pod. Pods consume node resources and PVCs consume PV resources.

If you intend to create Kubernetes persistent volumes, sufficient block volume quota must be available in each availability domain to meet the persistent volume claim. Persistent volume claims must request a minimum of 50 gigabytes

You can define and apply a persistent volume claim to your cluster, which in turn creates a persistent volume that's bound to the claim. A claim is a block storage volume in the underlying IaaS provider that's durable and offers persistent storage, enabling your data to remain intact, regardless of whether the containers that the storage is connected to are terminated.

With Oracle Cloud Infrastructure as the underlying IaaS provider, you can provision persistent volume claims by attaching volumes from the Block Storage service.





<https://oracle.github.io/weblog-kubernetes-operator/faq/oci-fss-pv/>
<https://kubernetes.io/docs/concepts/storage/persistent-volumes/>

Question: 4

You are a consumer of Oracle Cloud Infrastructure (OCI) Streaming service. Which API should you use to read and process the stream?

- A. ListMessages
- B. GetMessage
- C. GetObject
- D. ReadMessages

Answer: B

Explanation:

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

Building consumers to read and process messages from a stream using the [GetMessages](#) API .

How Streaming Works:

The Streaming service provides a robust, scalable mechanism that you can use to produce and consume high volumes of data between application components.

Here's how Streaming works: a producer publishes messages to a stream, which is an append-only log. These messages are distributed among the partitions using the message's key.

Streams are divided into a number of partitions for scalability. Partitions allow you to distribute a stream by splitting messages across multiple nodes (or brokers). Each partition can be placed on a separate machine to allow multiple consumers to read a stream in parallel. Multiple consumers can read from any partition regardless of where the partition is hosted.

A consumer can read messages from one or more streams. Each message within a stream is marked with an offset value, so a consumer can pick up where it left off if it is interrupted.

You can use the Streaming service by:

- A. Creating a stream using the Console or API.
- B. Using a producer to publish data to the stream.
- C. Building consumers to read and process messages from a stream using the [GetMessages](#) API . .

Reference:

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

Question: 5

Which two handle Oracle Functions authentication automatically?

- A. Oracle Cloud Infrastructure SDK
- B. cURL
- C. Oracle Cloud Infrastructure CLI
- D. Signed HTTP Request
- E. Fn Project CLI

Answer: CE

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.

Reference:

<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>