



- Expert Verified, Online, **Free**.

Which statement accurately describes Oracle Cloud Infrastructure (OCI) Load Balancer integration with OCI Container Engine for Kubernetes (OKE)? (Choose the best answer.)

- A. OKE service provisions an OCI Load Balancer instance for each Kubernetes service with LoadBalancer type in the YAML configuration.
- B. OCI Load Balancer instance provisioning is triggered by OCI Events service for each Kubernetes service with LoadBalancer type in the YAML configuration.
- C. OCI Load Balancer instance must be manually provisioned for each Kubernetes service that requires traffic balancing.
- D. OKE service provisions a single OCI Load Balancer instance shared with all the Kubernetes services with LoadBalancer type in the YAML configuration.

Suggested Answer: B

Reference:

<https://oracle.github.io/weblogic-kubernetes-operator/faq/oci-lb/>

  **n00pster** Highly Voted 4 years, 8 months ago

Its A

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

Load balancer services you create appear in the Console. However, do not use the Console (or the Oracle Cloud Infrastructure CLI or API) to modify load balancer services. Any modifications you make will either be reverted by Container Engine for Kubernetes or will conflict with its operation and possibly result in service interruption.

upvoted 7 times

  **Saidasa** Most Recent 3 years, 1 month ago

Namaste

I would vote for D

upvoted 1 times

  **fk_t_boss** 3 years, 3 months ago

has to be A

B: is wrong as the Events service is not in the context of deploying OKE services

C: is clearly wrong

D: is wrong, for the LB we can specify the shape so this statement can not be true "shared with all the Kubernetes services with LoadBalancer type"

upvoted 1 times

  **amitmishracs** 4 years, 1 month ago

which answer is correct D or B?

upvoted 1 times

  **issaprank** 4 years, 7 months ago

Either A or D.

upvoted 1 times

  **Sai19** 4 years, 7 months ago

It's D

upvoted 1 times

  **firstvishal** 4 years, 8 months ago

Correct answer is D

upvoted 1 times

  **wdahman** 4 years, 8 months ago

D no doubt

upvoted 1 times

🗨️ 👤 **FPM** 4 years, 1 month ago

D is wrong. Try it. No doubt. Correct answer is A.

upvoted 1 times

🗨️ 👤 **adesmaster** 4 years, 8 months ago

its's A OCI OKE creates one LB for each service in the yaml file. LB its different of ingress and LB doesn't have paths or routes so you can't share different OKE services with one LB.

upvoted 3 times

🗨️ 👤 **jsg** 4 years, 8 months ago

Ans is D

upvoted 2 times

🗨️ 👤 **Jatindra** 4 years, 8 months ago

Answer: D

upvoted 2 times

🗨️ 👤 **melb** 4 years, 8 months ago

It is D...

upvoted 1 times

🗨️ 👤 **DikshaGarg** 4 years, 8 months ago

i have checked on other sites. Answer is either A or B.

upvoted 1 times

🗨️ 👤 **rider** 4 years, 8 months ago

can you give details of that other websites?

upvoted 2 times

🗨️ 👤 **papayahead** 4 years, 8 months ago

I think it's C.

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

When you create a deployment, you can optionally create a load balancer service to distribute traffic between the nodes assigned to the deployment. The key fields in the configuration of a load balancer service are the type of service being created and the ports that the load balancer will listen to.

upvoted 1 times

🗨️ 👤 **papayahead** 4 years, 8 months ago

ok, I stand corrected... should be A.

B is incorrect because it's not triggered by OCI events service.

D is incorrect because each service gets a LB. in other words, you can't have one LB for all services.

C is incorrect because it's not manual, it's automatically created based on the LoadBalancer type configured in the yaml file.

upvoted 9 times

🗨️ 👤 **rider** 4 years, 8 months ago

Should be A. Correct

upvoted 4 times

🗨️ 👤 **DikshaGarg** 4 years, 8 months ago

i think its D

upvoted 2 times

Per CAP theorem, in which scenario do you NOT need to make any trade-off between the guarantees? (Choose the best answer.)

- A. when there are no network partitions
- B. when the system is running in the cloud
- C. when the system is running on-premise
- D. when you are using load balancers

Suggested Answer: A

Community vote distribution

A (100%)

  **gfhbox0083** Highly Voted 4 years, 8 months ago

The provided Answer is Correct.

The CAP theorem implies, that in the presence of a network partition, one has to choose between consistency and availability.

https://en.wikipedia.org/wiki/CAP_theorem#cite_note-4

upvoted 8 times



  **Saidasa** Most Recent 3 years, 1 month ago

Selected Answer: A

Namaste

A is correct

upvoted 1 times

  **Jatindra** 4 years, 8 months ago

Answer : A

upvoted 3 times

You have two microservices, A and B running in production. Service A relies on APIs from service B. You want to test changes to service A without deploying all of its dependencies, which includes service B.



Which approach should you take to test service A? (Choose the best answer.)

- A. Test against production APIs.
- B. Test using API mocks.
- C. There is no need to explicitly test APIs.
- D. Test the APIs in private environments.

Suggested Answer: B

Reference:

<https://docs.oracle.com/cloud/apiary/tools/mock-server/index.html>

  **gfhbox0083** Highly Voted 4 years, 8 months ago

The provided answer and link are correct.

Another use case via <https://docs.oracle.com/en/solutions/build-governance-app-oracle-paas/test-custom-apis.html#GUID-16C1880E-00ED-44E3-9767-A9C030D36653>


upvoted 5 times

  **Saidasa** Most Recent 3 years, 1 month ago

Namaste

B is correct

upvoted 1 times

  **Jatindra** 4 years, 8 months ago

Answer : B

upvoted 3 times

In a Linux environment, what is the default location of the configuration file that Oracle Cloud Infrastructure CLI uses for profile information?
(Choose the best answer.)

- A. /etc/.oci/config
- B. /usr/local/bin/config
- C. \$HOME/.oci/config
- D. /usr/bin/oci/config

Suggested Answer: C

Reference:

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

Community vote distribution

C (100%)

  **gfhbox0083** Highly Voted 4 years, 8 months ago

The provided answer is correct.

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

upvoted 5 times

  **Saidasa** Most Recent 3 years, 1 month ago

Selected Answer: C

Namaste



Answer is C i.e. ~/.oci/config

upvoted 1 times

  **firstvishal** 4 years, 8 months ago

Given answer is correct

upvoted 2 times

  **Jatindra** 4 years, 8 months ago

Answer : C

upvoted 1 times

Which statement is incorrect with regards to the Oracle Cloud Infrastructure (OCI) Notifications service? (Choose the best answer.)

- A. Notification topics may be assigned as the action performed by an OCI Events configuration.
- B. OCI Alarms can be configured to publish to a notification topic when triggered.
- C. An OCI function may subscribe to a notification topic.
- D. A subscription can forward notifications to an HTTPS endpoint.
- E. A subscription can integrate with PagerDuty events.
- F. It may be used to receive an email each time an OCI Autonomous Database backup is completed.

Suggested Answer: E

Reference:

<https://www.pagerduty.com/docs/guides/oracle-cloud-infrastructure/>

Community vote distribution

F (100%)

🗳️ 👤 **n00pster** Highly Voted 👍 4 years, 8 months ago

Probably F. Since all other answers are directly related to notification service. For F you need event service first.
upvoted 6 times

🗳️ 👤 **juanlopez2011** Most Recent 🕒 1 year, 7 months ago

Selected Answer: F

F is correct
upvoted 1 times

🗳️ 👤 **Saidasa** 3 years, 1 month ago

Selected Answer: F

Namaste

F is correct
upvoted 1 times

🗳️ 👤 **fk_t_boss** 3 years, 3 months ago

F is "best" answer as needs "event", In option A it explicitly declares that the event is configured. No such statement exists in option F. If you implicitly accept the existence of the event configuration for F all the options are correct. Questions like this should never appear in an exam where many candidates will have English as a second language. Options should all be clear and explicit in their intent. This is like a "trick" question.
upvoted 1 times

🗳️ 👤 **Paasu** 4 years, 4 months ago

F is the answer
upvoted 1 times

🗳️ 👤 **Ramesh88** 4 years, 7 months ago

F is the Answer
upvoted 1 times

🗳️ 👤 **Sai19** 4 years, 7 months ago

F is the correct answer
upvoted 1 times

🗳️ 👤 **cloudgrimm** 4 years, 7 months ago

All of them seem to be correct, F seems to be the answer
upvoted 1 times

🗳️ 👤 **av1pnupv2gsv** 4 years, 7 months ago

it should be F because Supported services include Compute, LB, Block & Object Storage, Notifications & Streaming but it doesn't include DBs
upvoted 3 times

🗨️ 👤 **firstvishal** 4 years, 8 months ago

F is the correct answer
upvoted 3 times

🗨️ 👤 **rider** 4 years, 8 months ago

it cannot be E. probably F is the answer.
upvoted 2 times

🗨️ 👤 **Jatindra** 4 years, 8 months ago

Answer : F
upvoted 4 times

🗨️ 👤 **SivaRamaKrishna** 4 years, 8 months ago

All the answers seems to be CORRECT statements. There is no INCORRECT statement. E is also a correct statement.

<https://www.ateam-oracle.com/oci-native-monitoring-and-alerting-for-autonomous-database-and-db-systems>

upvoted 3 times

🗨️ 👤 **ArshadRock** 4 years, 8 months ago

E seems to be correct answer
upvoted 2 times

With the volume of communication that can happen between different components in cloud-native applications, it is vital to not only test functionality, but also service resiliency.

Which statement is true with regards to service resiliency? (Choose the best answer.)

- A. Resiliency is about recovering from failures without downtime or data loss.
- B. A goal of resiliency is not to bring a service to a functioning state after a failure.
- C. Resiliency testing can be only done in a test environment.
- D. Resiliency is about avoiding failures.

Suggested Answer: D

Community vote distribution

A (100%)

🗨️ **examtaker1** Highly Voted 4 years, 8 months ago

I go with A.

Resiliency is the ability of your system to react to failure and still remain functional. It's not about avoiding failure, but accepting failure and constructing your cloud-native services to respond to it. You want to return to a fully functioning state quickly as possible.

<https://docs.microsoft.com/en-us/dotnet/architecture/cloud-native/resiliency>

upvoted 11 times

🗨️ **Saidasa** Most Recent 3 years, 1 month ago

Selected Answer: A

Namaste

A is correct being a basic definition of resiliency.

upvoted 1 times

🗨️ **gogoa** 3 years, 3 months ago

Selected Answer: A

A is correct, basic definition of resiliency

upvoted 1 times

🗨️ **Minghon** 3 years, 4 months ago

Selected Answer: A

Resiliency is the ability of your system to react to failure and still remain functional. It's not about avoiding failure, but accepting failure and constructing your cloud-native services to respond to it.

Reference: <https://docs.microsoft.com/en-us/dotnet/architecture/cloud-native/resiliency>

upvoted 2 times

🗨️ **Janin** 4 years, 6 months ago

A. Resiliency is about recovering from failures without downtime or data loss

upvoted 1 times

🗨️ **Adijobi** 4 years, 7 months ago

It's A this is basic SRE knowledge. The idea is that failures are detected so for example a service node fails synthetic monitoring or alive checks. Kubernetes stops traffic to the node at the load balancer. Then kubernetes rebuilds the node and replaces the defective one in the cluster.

Another real world application is ChaosMonkey which intentionally creates faults in prod (big balls energy) and tests that the system is resilient enough to tolerate this without actual downtime or data loss.

upvoted 1 times

🗨️ **rider** 4 years, 8 months ago

answer is A

upvoted 1 times

🗨️ **Ani_1811** 4 years, 8 months ago

We are continuing
despite failures so we are not avoiding failure here i go with A
upvoted 3 times

🗨️ 👤 **Jatindra** 4 years, 8 months ago

Answer : D

upvoted 1 times

🗨️ 👤 **xteuaiporhqdryvabm** 4 years, 8 months ago

It is A .

Resiliency and availability refers to the ability of a system to continue operating, despite the failure or sub-optimal performance of some of its components.

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

upvoted 2 times

🗨️ 👤 **goodh32** 4 years, 8 months ago

Correct Answer is is A

upvoted 3 times

🗨️ 👤 **goodh32** 4 years, 8 months ago

Resiliency is the ability to recover from failures and continue to function. It isn't about avoiding failures but accepting the fact that failures will happen and responding to them in a way that avoids downtime or data loss

<https://docs.microsoft.com/en-us/dotnet/architecture/microservices/implement-resilient-applications/#:~:text=Resiliency%20is%20the%20ability%20to,avoids%20downtime%20or%20data%20loss.>

upvoted 4 times

🗨️ 👤 **DikshaGarg** 4 years, 8 months ago

Answer is correct.. Its D

upvoted 2 times

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

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

Suggested Answer: BE

Reference:

<https://www.oracle.com/webfolder/technetwork/tutorials/obe/oci/oke-full/index.html>

Community vote distribution

BE (100%)

🗨️ 👤 **Saidasa** 3 years, 1 month ago

Selected Answer: BE

Namaste

B & E are correct

upvoted 1 times

🗨️ 👤 **Minghon** 3 years, 4 months ago

Selected Answer: BE

To access the cluster using a local installation of kubectl:

1. Generate an API signing key pair (if you don't already have one).
2. Upload the public key of the API signing key pair.
3. Install and configure the Oracle Cloud Infrastructure CLI.
4. Set up the kubeconfig file.

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

upvoted 2 times

🗨️ 👤 **Jatindra** 4 years, 8 months ago

Answer : B,E

upvoted 2 times

🗨️ 👤 **gfhbox0083** 4 years, 8 months ago

B, and E for sure.

To access the cluster using a local installation of kubectl and the Kubernetes Dashboard:

Generate an API signing key pair (if you don't already have one).

Upload the public key of the API signing key pair.

Install and configure the Oracle Cloud Infrastructure CLI.

Set up the kubeconfig file.

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

upvoted 2 times

You have a containerized app that requires an Autonomous Transaction Processing (ATP) Database. Which option is not valid for connecting to ATP from a container in Kubernetes? (Choose the best answer.)

- A. Enable Oracle REST Data Services for the required schemas and connect via HTTPS.
- B. Create a Kubernetes secret with contents from the instance Wallet files. Use this secret to create a volume mounted to the appropriate path in the application deployment manifest.
- C. Use Kubernetes secrets to configure environment variables on the container with ATP instance OCID, and OCI API credentials. Then use the CreateConnection API endpoint from the service runtime.
- D. Install the Oracle Cloud Infrastructure Service Broker on the Kubernetes cluster and deploy ServiceInstance and ServiceBinding resources for ATP. Then use the specified binding name as a volume in the application deployment manifest.

Suggested Answer: D

Reference:

<https://blogs.oracle.com/developers/creating-an-atp-instance-with-the-oci-service-broker>

Community vote distribution

A (100%)

 **gfhbox0083** Highly Voted 4 years, 8 months ago

The Question is about Which option is Not valid for connecting to ATP from a container in Kubernetes?

The provided answers and link are valid in this case. I believe B should be the answer as an Invalid option
upvoted 6 times

 **Saidasa** Most Recent 3 years, 1 month ago

Selected Answer: A

Namaste


A is incorrect. Rest are correct.

upvoted 1 times

 **fk_t_boss** 3 years, 3 months ago

I think is is C: if you think otherwise please post a link to the oci "CreateConnection API", I dont think it exists

upvoted 1 times

 **shivakkr** 3 years, 10 months ago


Ans is A

upvoted 2 times

 **firstvishal** 4 years, 7 months ago

Correct answer is A

upvoted 1 times


 **mbanzai** 4 years, 8 months ago

Answer : A

B,C,D are valid as they legit steps to Creating An ATP Instance With The OCI Service Broker, described in the link

<https://blogs.oracle.com/developers/creating-an-atp-instance-with-the-oci-service-broker>

upvoted 3 times

 **jsg** 4 years, 8 months ago


Ans is A

upvoted 1 times

 **Jatindra** 4 years, 8 months ago

Answer : A



upvoted 1 times

 **n00pster** 4 years, 8 months ago

Either B or C. Looking at link below, I think its C.

<https://blogs.oracle.com/developers/oracle-functions-connecting-to-an-atp-database-with-a-wallet-stored-as-secrets>

upvoted 4 times

  **papayahead** 4 years, 7 months ago

I agree, C is the correct answer.

upvoted 2 times

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? (Choose the best answer.)

- A. OCI Compute
- B. OCI Container Engine for Kubernetes
- C. OCI Resource Manager
- D. OCI API Gateway

Suggested Answer: C

  **gfhbox0083** Highly Voted 4 years, 8 months ago

C, for sure



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.

upvoted 6 times

  **c_tac** Most Recent 1 year, 6 months ago



C is correct

upvoted 1 times

  **gogoa** 3 years, 3 months ago

c ANSWER

upvoted 2 times

  **Jatindra** 4 years, 8 months ago

Answer : C

upvoted 3 times

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? (Choose the best answer.)

- A. The request must include an authorization signing string including (but not limited to) x-content-sha256, content-type, and content-length headers.
- B. The Content-Type header must be set to application/json
- C. An HTTP 401 will be returned if the client's clock is skewed more than 5 minutes from the server's.
- D. The request does not require an Authorization header.

Suggested Answer: D

Reference:

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

Community vote distribution

D (100%)

 **mbanzai** Highly Voted 4 years, 8 months ago

D

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

and the Authorization header.

upvoted 6 times


 **Saidasa** Most Recent 3 years, 1 month ago

Selected Answer: D

Namaste

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

upvoted 1 times

 **Osong** 4 years, 7 months ago

D

To see why B and C are correct. See <https://docs.cloud.oracle.com/en-us/iaas/Content/API/Concepts/usingapi.htm>

A is also correct, while if you are missing the header a 401 Unauthorized response will be sent to the client. See

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

You know an answer is from a bot if it cites no source nor argument

upvoted 2 times

 **firstvishal** 4 years, 7 months ago

I think answer is D- I got practice exams from Udemy and in that answer is A, but when I read the article at oracle I found D must be the right answer:-

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.
See the remaining sections in this topic for details about these steps.

upvoted 3 times

  **Jatindra** 4 years, 8 months ago

Answer : A

upvoted 1 times

  **juanlopez2011** 1 year, 7 months ago

No, is the option D. The questions is for the incorrect

upvoted 1 times

  **gfhbox0083** 4 years, 8 months ago

D, for sure.

The request 'does' require an Authorization header.

More details via the provided link

upvoted 3 times


You are working on a serverless DevSecOps application using Oracle Functions. You have deployed a Python function that uses the Oracle Cloud Infrastructure (OCI) Python SDK to stop any OCI Compute instance that does not comply with your corporate security standards. There are 3 non-compliant OCI Compute instances.

However, when you invoke this function none of the instances were stopped.

How should you troubleshoot this? (Choose the best answer.)

- A. There is no way to troubleshoot a function running on Oracle Functions.
- B. Enable function logging in the OCI console, include some print statements in your function code and use logs to troubleshoot this.
- C. Enable function remote debugging in the OCI console, and use your favorite IDE to inspect the function running on Oracle Functions.
- D. Enable function tracing in the OCI console, and go to OCI Monitoring console to see the function stack trace.

Suggested Answer: C

 **gfhbox0083** Highly Voted 4 years, 8 months ago

Would go with B, for this time.

Debugging would be in case of an error returned by the function.

more info via

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


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

upvoted 9 times

 **thulasi39** Most Recent 4 years, 7 months ago

Ans: B

upvoted 3 times

 **Adijobi** 4 years, 7 months ago


Gosh after doing research I have to say B actually looks correct. Fn's own troubleshooting docs seem to lean this way ---JANKY!!!---

upvoted 1 times

 **firstvishal** 4 years, 7 months ago

Correct answer is B

upvoted 1 times

 **mbanzai** 4 years, 8 months ago

B

There is no evidence that the other system works.

upvoted 2 times

Which is NOT a valid option to execute a function deployed on Oracle Functions? (Choose the best answer.)

- A. Send a signed HTTP requests to the function's invoke endpoint
- B. Invoke from Oracle Cloud Infrastructure CLI
- C. Invoke from Docker CLI
- D. Trigger by an event in Oracle Cloud Infrastructure Events service
- E. Invoke from Fn Project CLI

Suggested Answer: C

Reference:

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

🗨️ 👤 **Irgond07** 3 years, 2 months ago

Please read the question carefully.. i would go with answer D

upvoted 1 times

🗨️ 👤 **Jatindra** 4 years, 8 months ago

Answer : C

upvoted 3 times

🗨️ 👤 **gfhbox0083** 4 years, 8 months ago

C, for sure.

You can invoke a function that you've deployed to Oracle Functions in different ways:

Using the Fn Project CLI.

Using the Oracle Cloud Infrastructure CLI.

Using the Oracle Cloud Infrastructure SDKs.

Making a signed HTTP request to the function's invoke endpoint. Every function has an invoke endpoint.

Other Oracle Cloud services (for example, triggered by an event in the Events service) or from external services.

upvoted 4 times

You are developing a polyglot serverless application using Oracle Functions.
Which language cannot be used to write your function code? (Choose the best answer.)

- A. PL/SQL
- B. Python
- C. Node.js
- D. Go
- E. Java

Suggested Answer: A



Reference:

https://static.rainfocus.com/oracle/oow19/sess/1550723907257001SAfg/PF/Serverless%20-%20A%20Game%20Changer%20in%20Cloud%20Computing%20%281%29_1568730727062001lhkS.pdf

(18)

Community vote distribution


A (100%)

  **gfhbox0083** Highly Voted 4 years, 8 months ago

A, for sure.

Oracle Functions is based on open source Fn Project which provides function development kits (FDKs) for Java (including GraalVM native image), Python, Node, Go and Ruby, and in addition, lets you bring your own Dockerfile.

upvoted 5 times

  **ProbaMarci** Most Recent 1 year, 7 months ago

Selected Answer: A

A. PL/SQL is used for programming in database like writing triggers or stores procedures. Functions have nothing to do with that, it executes on a virtual machine (in a container) and not in a database.

upvoted 1 times

  **Saidasa** 3 years, 1 month ago

Namaste



Answer is A

upvoted 1 times

  **RobChef** 3 years, 9 months ago

PL/SQL is correct answer

upvoted 1 times

  **vpexam** 4 years, 8 months ago

node.js, go, python and java all can be used except PL/ SQL

upvoted 1 times



Which two statements accurately describe an Oracle Functions application? (Choose two.)

- 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

Suggested Answer: AB

Reference:

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

  **shubhachandra** Highly Voted 4 years, 8 months ago

Answer should be DE

Applications



In Oracle Functions, an application is:

a logical grouping of functions

a common context to store configuration variables that are available to all functions in the application

a way to ensure function runtime isolation

upvoted 14 times

  **gfobox0083** Highly Voted 4 years, 8 months ago

D, E for sure.

The provided link is correct. It contains the correct answers.

In Oracle Functions, an application is:

A logical grouping of functions

A common context to store configuration variables that are available to all functions in the application.

A way to ensure function runtime isolation.



upvoted 8 times

  **Saidasa** Most Recent 3 years, 1 month ago

Namaste

D, E

upvoted 1 times

  **sranje** 3 years, 3 months ago

D&E is true

upvoted 1 times

  **RobChef** 3 years, 9 months ago

ANS - DE

Applications

In Oracle Functions, an application is:

a logical grouping of functions

a way to allocate and configure resources for all functions in the application

a common context to store configuration variables that are available to all functions in the application

a way to ensure function runtime isolation

upvoted 1 times

  **RobChef** 3 years, 9 months ago

ANS - D,E

upvoted 1 times

🗨️ 👤 **Adijobi** 4 years, 7 months ago

fn functions are not docker images though fn can technically be deployed to docker such things are irrelevant in OCI
upvoted 1 times

🗨️ 👤 **firstvishal** 4 years, 7 months ago

D, E are the correct options
upvoted 2 times

🗨️ 👤 **rider** 4 years, 8 months ago

D and E as mentioned in video tutorials
upvoted 2 times

🗨️ 👤 **papayahead** 4 years, 8 months ago

Functions

In Oracle Functions, functions are:

small but powerful blocks of code that generally do one simple thing
grouped into applications

stored as Docker images in a specified Docker registry

invoked in response to a CLI command or signed HTTP request

upvoted 1 times

🗨️ 👤 **papayahead** 4 years, 8 months ago

The question is so tricky...

upvoted 1 times

🗨️ 👤 **papayahead** 4 years, 8 months ago

It asks about "Application", not "function", so I agree with D&E...

upvoted 2 times

🗨️ 👤 **t_kaite** 4 years, 8 months ago

I agree, the Answer should be D & E as referenced in the link provided under the answers.

upvoted 4 times

You are processing millions of files in an Oracle Cloud Infrastructure (OCI) Object Storage bucket. Each time a new file is created, you want to send an email to the customer and create an order in a database. The solution should perform and minimize cost. Which action should you use to trigger this email? (Choose the best answer.)

- A. Schedule a cron job that monitors the OCI Object Storage bucket and emails the customer when a new file is created.
- B. Use OCI Events service and OCI Notification service to send an email each time a file is created.
- C. Schedule an Oracle Function that checks the OCI Object Storage bucket every minute and emails the customer when a file is found.
- D. Schedule an Oracle Function that checks the OCI Object Storage bucket every second and emails the customer when a file is found.

Suggested Answer: C

  **gfhbox0083** Highly Voted 4 years, 8 months ago

Would go with B.

Use Notifications to get notified when Event Rules are triggered or alarms are breached, or to directly publish a message.



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

upvoted 11 times

  **sranje** Most Recent 3 years, 3 months ago



B is the correct answer

upvoted 1 times

  **gogoa** 3 years, 3 months ago


B, event driven

upvoted 1 times

  **scab** 4 years, 5 months ago


To minimize cost, i think C is better than B. what do you think ?

upvoted 1 times

  **tohegajaf** 3 years, 6 months ago



with function you pay per execution time. doesn't make sense to run it in a cron

upvoted 1 times

  **scab** 4 years, 5 months ago

To minimize cost, i think B is beter than B. what do you think ?

upvoted 1 times

  **tinojobi** 4 years, 7 months ago



B is the correct answer

upvoted 2 times

  **firstvishal** 4 years, 7 months ago


Correct ans is B

upvoted 2 times

  **vpexam** 4 years, 8 months ago

B is most relevant

upvoted 2 times

  **Chamstombs** 4 years, 8 months ago

The answer is B

upvoted 3 times

  **papayahead** 4 years, 8 months ago

I think the answer should be B

upvoted 4 times

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? (Choose the best answer.)

- A. 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.
- B. Create an OCI Notifications 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.
- C. 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.
- D. 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.

Suggested Answer: D

ddmoto Highly Voted 4 years, 8 months ago

Should be B

upvoted 12 times

devops_sams Most Recent 3 years, 1 month ago

Yes B is right

upvoted 1 times

Balazs_Molnar 3 years, 11 months ago

Probably B since there is no sign of email delivery service possibilities in <https://docs.cloud.oracle.com/en-us/iaas/Content/Notification/Concepts/notificationoverview.htm>

upvoted 1 times

Adijobi 4 years, 7 months ago

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

Yep B!!

upvoted 1 times

firstvishal 4 years, 7 months ago

Correct answer is B

upvoted 1 times

mbanzai 4 years, 8 months ago

B

You can only deliver events to certain Oracle Cloud Infrastructure services with a rule. Use the following services to create actions:

Notifications

Streaming

Functions

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

upvoted 3 times

A service you are deploying to Oracle Cloud Infrastructure (OCI) Container Engine for Kubernetes (OKE) uses a docker image from a private repository in OCI Registry (OCIR).

Which configuration is necessary to provide access to this repository from OKE? (Choose the best answer.)

- A. Add a generic secret on the cluster containing your identity credentials. Then specify a registryCredentials property in the deployment manifest.
- B. Create a docker-registry secret for OCIR with API key credentials on the cluster, and specify the imagePullSecret property in the application deployment manifest.
- C. Create a docker-registry secret for OCIR with identity Auth Token on the cluster, and specify the imagePullSecret property in the application deployment manifest.
- D. Create a dynamic group for nodes in the cluster, and a policy that allows the dynamic group to read repositories in the same compartment.

Suggested Answer: C

Reference:

<https://www.oracle.com/webfolder/technetwork/tutorials/obe/oci/oke-and-registry/index.html>

 **gfhbox0083** Highly Voted 4 years, 8 months ago

The provided answer C, is correct.

Additional info, via

https://oracle.github.io/learning-library/oci-library/DevOps/Container_Registry/Container_Registry_HOL.html

upvoted 6 times

 **Saidasa** Most Recent 3 years, 1 month ago

Namaste

C is the Answer

upvoted 1 times

Given a service deployed on Oracle Cloud Infrastructure Container Engine for Kubernetes (OKE), which annotation should you add in the sample manifest file below to specify a 400 Mbps load balancer? (Choose the best answer.)


```
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-kind: 400Mbps
- B. service.beta.kubernetes.io/oci-load-balancer-value: 400Mbps
- C. service.beta.kubernetes.io/oci-load-balancer-shape: 400Mbps
- D. service.beta.kubernetes.io/oci-load-balancer-size: 400Mbps

Suggested Answer: C

Reference:

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

 **gfhbox0083** Highly Voted 4 years, 8 months ago

C, for sure.

The shape of an Oracle Cloud Infrastructure load balancer specifies its maximum total bandwidth (that is, ingress plus egress). By default, load balancers are created with a shape of 100Mbps. Other shapes are available, including 400Mbps and 8000Mbps.

To specify an alternative shape for a load balancer, add the following annotation in the metadata section of the manifest file:

```
service.beta.kubernetes.io/oci-load-balancer-shape: <value>
```

upvoted 9 times

 **Saidasa** Most Recent 3 years, 1 month ago

Namaste

C is correct

upvoted 1 times

 **Balazs_Molnar** 3 years, 11 months ago

C there is no kind or value or size parameters just shape

upvoted 2 times


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.)


- A. 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'
- B. Set up the following dynamic group for your function's OCID: Name: read-file-dg Rule: resource.id = "'ocid1.fnfunc.oc1.phx.aaaaaaaakeobctakezjz5i4ujj7g25q7sx5mvr55pms6f4da'
- C. 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'
- D. 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'
- E. No policies are needed. By default, every function has read access to Object Storage buckets in the tenancy.


Suggested Answer: AC

 **ddmoto** Highly Voted 4 years, 8 months ago
should be a and B
upvoted 9 times


 **Saidasa** Most Recent 3 years, 1 month ago
Namaste

A and B
upvoted 1 times

 **firstvishal** 4 years, 7 months ago
Correct answers are A and B
upvoted 4 times

 **mbanzai** 4 years, 8 months ago
the required steps are:

- create a Dynamic Group with a rule to include Functions in a Compartment (or even functions with a specific name)
 - create a policy to grant this Dynamic Group the permission to read objects from the Bucket
- upvoted 1 times

 **melb** 4 years, 8 months ago
perfect agree A, B..first create dynamic group 'read-file-dg' and then add this function OCID to 'read-file-dg' as per the resource-principals creation rule.
upvoted 3 times

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? (Choose the best answer.)

- A. `kubectl logs pod nginx`
- B. `kubectl describe pod nginx`
- C. `kubectl get pod nginx`
- D. Through the Oracle Cloud Infrastructure Console

Suggested Answer: B

Reference:

<https://github.com/kubernetes/kubernetes/issues/49314>

 **gfhbox0083** Highly Voted 4 years, 8 months ago

B, is correct.

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.

<https://kubernetes.io/docs/tasks/debug-application-cluster/debug-application/#my-pod-stays-pending>

upvoted 5 times

 **examsycia** Most Recent 3 years, 2 months ago

I found this:

<https://www.oracle.com/cloud-native/container-engine-kubernetes/faq/>


the easiest way to troubleshoot the pod is to get the logs from the pod with the following commands:

- * Use `kubectl get pods` to get a list of running pods in your cluster.
- * From that list, copy the pod ID for the Oracle Cloud Infrastructure Service Broker.
- * Use `kubectl logs` to get the logs from the pod.

I have doubts because "kubectl logs pod nginx" is incorrect it'd be "kubectl logs <Pod_Name>" but it's Oracle's recommendation.

Alternative, B (gfhbox0083 is right)

upvoted 1 times

 **D_nice** 4 years, 8 months ago

Ans: B

upvoted 3 times

A pod security policy (PSP) is implemented in your Oracle Cloud Infrastructure Container Engine for Kubernetes cluster. Which rule can you use to prevent a container from running as root using PSP? (Choose the best answer.)

- A. NoPrivilege
- B. RunOnlyAsUser
- C. MustRunAsNonRoot
- D. forbiddenRoot

Suggested Answer: C

Reference:

<https://docs.bitnami.com/tutorials/secure-kubernetes-cluster-bsp/>

Community vote distribution

C (100%)

🗨️ 👤 **Saidasa** 3 years, 1 month ago

Selected Answer: C

Namaste

C is correct

upvoted 1 times

🗨️ 👤 **Ahmed_Shawky** 4 years, 7 months ago

yes C is The right answer

upvoted 1 times

🗨️ 👤 **gfhbox0083** 4 years, 8 months ago

C, for sure

Require the container to run without root privileges, rule: 'MustRunAsNonRoot'

<https://kubernetes.io/docs/concepts/policy/pod-security-policy/>

upvoted 4 times


What is one of the differences between a microservice and a serverless function? (Choose the best answer.)

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

Suggested Answer: B

Community vote distribution

A (100%)

 **gfhbox0083** Highly Voted 4 years, 8 months ago

I would go with A.

Microservices are best suited for long-running, complex applications that have significant resource and management requirements.


On the other hand, serverless functions only execute when needed. Once the execution is over, the computing instance that runs the code decommissions itself.

D is wrong, because both can be triggered by events, both are event driven

C is wrong, because Microservices can be stateless or stateful

B is wrong, because serverless functions can use datastore to store data.

upvoted 11 times

 **DikshaGarg** Highly Voted 4 years, 8 months ago

Answer is A.

upvoted 9 times

 **Saidasa** Most Recent 3 years, 1 month ago

Selected Answer: A

Namaste

A is the correct answer

upvoted 1 times

 **gogoa** 3 years, 3 months ago

A answer

upvoted 2 times

 **onlygrass** 4 years, 3 months ago

The correct answer should be B.

A basic principle of microservices is that each service manages its own data. Two services should not share a data store. Instead, each service is responsible for its own private data store, which other services cannot access directly.

Ref:

[https://docs.microsoft.com/en-us/azure/architecture/microservices/design/data-](https://docs.microsoft.com/en-us/azure/architecture/microservices/design/data-considerations#:~:text=A%20basic%20principle%20of%20microservices,other%20services%20cannot%20access%20directly.)


[considerations#:~:text=A%20basic%20principle%20of%20microservices,other%20services%20cannot%20access%20directly.](https://docs.microsoft.com/en-us/azure/architecture/microservices/design/data-considerations#:~:text=A%20basic%20principle%20of%20microservices,other%20services%20cannot%20access%20directly.)

Serverless applications have no place to store persistent data or files. They don't have a built-in database or permanent file system.

[https://www.oclc.org/developer/news/2019/storing-data-in-a-serverless-](https://www.oclc.org/developer/news/2019/storing-data-in-a-serverless-application.en.html#:~:text=Serverless%20applications%20have%20no%20place,export%20via%20a%20web%20application.)

[application.en.html#:~:text=Serverless%20applications%20have%20no%20place,export%20via%20a%20web%20application.](https://www.oclc.org/developer/news/2019/storing-data-in-a-serverless-application.en.html#:~:text=Serverless%20applications%20have%20no%20place,export%20via%20a%20web%20application.)

upvoted 2 times

 **Adijobi** 4 years, 7 months ago

Answer is A. I swear all the braindumps on the net must come from like two people because I've seen this wrong answer many times.

upvoted 2 times

 **firstvishal** 4 years, 7 months ago

Correct answer is A

upvoted 1 times


Which two "Action Type" options are NOT available in an Oracle Cloud Infrastructure (OCI) Events rule definition? (Choose two.)

- A. Notifications
- B. Functions
- C. Streaming
- D. Email
- E. Slack

Suggested Answer: AC

Reference:

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

 **WasimMuh** Highly Voted 4 years, 8 months ago

i think it's D and E

upvoted 8 times


 **gfhbox0083** Highly Voted 4 years, 8 months ago

D, E for sure.

The Events service can invoke any of the following services by delivering an event message for processing: Notifications, Streaming, Functions.

Although the provided link is correct

upvoted 7 times

 **Balazs_Molnar** Most Recent 3 years, 11 months ago

D, E is correct do not mismatch it with Notifications where A,C would be the correct, but here we are talking about Events and not Notifications.

upvoted 1 times

 **HKMAMUN** 4 years, 7 months ago

According to the provided reference link (<https://docs.cloud.oracle.com/en-us/iaas/Content/Events/Task/managingrules.htm>) the answers is D & E

upvoted 1 times

 **Ahmed_Shawky** 4 years, 7 months ago

D, E for sure.

upvoted 1 times

 **firstvishal** 4 years, 7 months ago

Slack and Email so answer must be D, E

Prerequisites for Creating Rules

Action resources: You must have resources already set up to specify as an action. The Events service invokes the action specified in the rule by delivering the event message to action resources, which can include topics, streams, or functions. Every rule must have at least one action. The Events service can invoke any of the following services by delivering an event message for processing:

Notifications

Streaming

Functions

upvoted 2 times

 **firstvishal** 4 years, 7 months ago

Slack and Email so answer must be D, E

upvoted 1 times

 **DikshaGarg** 4 years, 8 months ago

answer is D & E

upvoted 5 times







Which is NOT a supported SDK on Oracle Cloud Infrastructure (OCI)? (Choose the best answer.)

- A. Go SDK
- B. Java SDK
- C. .NET SDK
- D. Ruby SDK
- E. Python SDK

Suggested Answer: C

Reference:

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

-  **JasbirSohi** 3 years, 3 months ago
as per latest mentioned SDK's are now supported
upvoted 1 times
-  **vpexam** 4 years, 8 months ago
as per the training videos, there is no mention of .NET
upvoted 1 times
-  **D_nice** 4 years, 8 months ago
C is correct
upvoted 2 times
-  **mbanzai** 4 years, 8 months ago
recently .net has become supported, wondering if this questions are gonna be part of the exam
upvoted 4 times
-  **exam67** 3 years, 3 months ago
confirm that C used to be correct. It is not anymore. ".NET" is supported
upvoted 1 times
-  **DikshaGarg** 4 years, 8 months ago
Answer is correct. It's C.
upvoted 4 times


You want to push a new image in the Oracle Cloud Infrastructure (OCI) Registry.
Which two actions do you need to perform? (Choose two.)

- A. Assign a tag via Docker CLI to the image.
- B. Generate an auth token to complete the authentication via Docker CLI.
- C. Generate an API signing key to complete the authentication via Docker CLI.
- D. Assign an OCI defined tag via OCI CLI to the image.
- E. Generate an OCI tag namespace in your repository.

Suggested Answer: AB

Reference:

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

 **gfhbox0083** Highly Voted 4 years, 8 months ago

B, and A for sure.

To push an Image to Oracle Cloud Infrastructure Registry

1. Create an auth token for use with Oracle Cloud Infrastructure Registry
2. Log in to Oracle Cloud Infrastructure Registry from the Docker CLI
3. Pull a test image from DockerHub
4. Tag the image
5. Push the image to Oracle Cloud Infrastructure Registry using the Docker CLI
6. Verify the image has been pushed to Oracle Cloud Infrastructure Registry using the Console

<https://www.oracle.com/webfolder/technetwork/tutorials/obe/oci/registry/index.html>


upvoted 7 times


How can you find details of the tolerations field for the sample YAML file below? (Choose the best answer.)


```
apiVersion: v1
kind: Pod
metadata:
  name: busybox
  namespace: default
spec:
  containers:
  - image: busybox
    command:
    - sleep
    - "3600"
    imagePullPolicy: IfNotPresent
    name: busybox
  restartPolicy: Always
  tolerations:
  ...
```

- A. kubectl list pod.spec.tolerations
- B. kubectl explain pod.spec.tolerations
- C. kubectl describe pod.spec.tolerations
- D. kubectl get pod.spec.tolerations


Suggested Answer: D


-  **hopewins** Highly Voted 4 years, 8 months ago


The answer is B : kubectl explain pod.spec.tolerations
 Kubectl explain lists the files for supported resources.
<https://kubernetes.io/docs/reference/generated/kubectl/kubectl-commands#explain>
<https://stackoverflow.com/questions/58790754/using-the-kubectl-explain-command>
 upvoted 9 times
-  **surendraselokar** Most Recent 3 years, 10 months ago


Correct answer is B
 upvoted 2 times
-  **s_muyo** 4 years, 6 months ago


The correct answer is B.

https://jamesdefabia.github.io/docs/user-guide/kubectl/kubectl_explain/
 upvoted 2 times
-  **thulasi39** 4 years, 7 months ago

Correct answer is B.
 upvoted 1 times
-  **firstvishal** 4 years, 7 months ago

Correct answer is B
 upvoted 1 times
-  **mbanzai** 4 years, 8 months ago

Answer B.
 kubectl explain pod.spec.tolerations is the only one that actually works!
 upvoted 3 times
-  **n00pster** 4 years, 8 months ago

I tried them too.. agree only B worked! :)
 upvoted 2 times
-  **ddmoto** 4 years, 8 months ago

B is the answer

upvoted 3 times

  **DikshaGarg** 4 years, 8 months ago

yes , Answer is C.

upvoted 1 times

  **DikshaGarg** 4 years, 8 months ago

Answer is B, <https://itnext.io/understanding-kubectl-explain-9d703396cc8>

upvoted 5 times

  **gfhbox0083** 4 years, 8 months ago

I will go with C.

We can retrieve a lot more information about each of pods using "kubectl describe pod".

<https://kubernetes.io/docs/tasks/debug-application-cluster/debug-application-introspection/>

upvoted 1 times