



VCE & PDF

Pass4itSure.com

<https://www.pass4itsure.com/associate-cloud-engineer.html>
2024 Latest pass4itsure ASSOCIATE-CLOUD-ENGINEER PDF and VCE
dumps Download

ASSOCIATE-CLOUD-ENGINEER^{Q&As}

Associate Cloud Engineer

**Pass Google ASSOCIATE-CLOUD-ENGINEER Exam
with 100% Guarantee**

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

<https://www.pass4itsure.com/associate-cloud-engineer.html>

100% Passing Guarantee
100% Money Back Assurance

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

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





QUESTION 1

You have a Compute Engine instance hosting a production application. You want to receive an email if the instance consumes more than 90% of its CPU resources for more than 15 minutes. You want to use Google services. What should you do?

A. 1. Create a consumer Gmail account.

2.

Write a script that monitors the CPU usage.

3.

When the CPU usage exceeds the threshold, have that script send an email using the Gmail account and smtp.gmail.com on port 25 as SMTP server.

B. 1. Create a Stackdriver Workspace, and associate your Google Cloud Platform (GCP) project with it.

2.

Create an Alerting Policy in Stackdriver that uses the threshold as a trigger condition.

3.

Configure your email address in the notification channel.

C. 1. Create a Stackdriver Workspace, and associate your GCP project with it.

2.

Write a script that monitors the CPU usage and sends it as a custom metric to Stackdriver.

3.

Create an uptime check for the instance in Stackdriver.

D. 1. In Stackdriver Logging, create a logs-based metric to extract the CPU usage by using this regular expression: CPU Usage: ([0-9] {1,3})%

2.

In Stackdriver Monitoring, create an Alerting Policy based on this metric.

3.

Configure your email address in the notification channel.

Correct Answer: B

QUESTION 2

Your company wants to migrate their on-premises workloads to Google Cloud. The current on-premises workloads



consist of:

1.
A Flask web application
2.
A backend API
3.
A scheduled long-running background job for ETL and reporting

You need to keep operational costs low. You want to follow Google-recommended practices to migrate these workloads to serverless solutions on Google Cloud. What should you do?

- A. Migrate the web application to App Engine and the backend API to Cloud Run. Use Cloud Tasks to run your background job on Compute Engine.
- B. Migrate the web application to App Engine and the backend API to Cloud Run. Use Cloud Tasks to run your background job on Cloud Run.
- C. Run the web application on a Cloud Storage bucket and the backend API on Cloud Run. Use Cloud Tasks to run your background job on Cloud Run.
- D. Run the web application on a Cloud Storage bucket and the backend API on Cloud Run. Use Cloud Tasks to run your background job on Compute Engine.

Correct Answer: B

<https://cloud.google.com/architecture/migration-to-gcp-deploying-your-workloads>

QUESTION 3

You created a Kubernetes deployment by running `kubectl run nginx image=nginx labels=app=prod`. Your Kubernetes cluster is also used by a number of other deployments. How can you find the identifier of the pods for this nginx deployment?

- A. `kubectl get deployments ?output=pods`
- B. `gcloud get pods ?selector="app=prod"`
- C. `kubectl get pods -l "app=prod"`
- D. `gcloud list gke-deployments -filter={pod }`

Correct Answer: C

Explanation: This command correctly lists pods that have the label `app=prod`. When creating the deployment, we used the label `app=prod` so listing pods that have this label retrieve the pods belonging to nginx deployments. You can list pods by using Kubernetes CLI `kubectl get pods`.

Ref: <https://kubernetes.io/docs/tasks/access-application-cluster/list-all-running-container-images/> Ref: <https://kubernetes.io/docs/tasks/access-application-cluster/list-all-running-container-images/#list-containers-filtering-by->



pod-label

QUESTION 4

You created a cluster.YAML file containing

```
resources: name: cluster type: container.v1.cluster properties: zone: europe-west1-b cluster: description: My GCP ACE cluster initialNodeCount: 2
```

You want to use Cloud Deployment Manager to create this cluster in GKE. What should you do?

- A. `gcloud deployment-manager deployments create my-gcp-ace-cluster --config cluster.yaml`
- B. `gcloud deployment-manager deployments create my-gcp-ace-cluster --type container.v1.cluster -- config cluster.yaml`
- C. `gcloud deployment-manager deployments apply my-gcp-ace-cluster --type container.v1.cluster --config cluster.yaml`
- D. `gcloud deployment-manager deployments apply my-gcp-ace-cluster --config cluster.yaml`

Correct Answer: D

Explanation: `gcloud deployment-manager deployments create` creates deployments based on the configuration file. (Infrastructure as code). All the configuration related to the artifacts is in the configuration file. This command correctly creates a cluster based on the provided `cluster.yaml` configuration file.

Ref: <https://cloud.google.com/sdk/gcloud/reference/deployment-manager/deployments/create>

QUESTION 5

You have a developer laptop with the Cloud SDK installed on Ubuntu. The Cloud SDK was installed from the Google Cloud Ubuntu package repository. You want to test your application locally on your laptop with Cloud Datastore. What should you do?

- A. Export Cloud Datastore data using `gcloud datastore export`.
- B. Create a Cloud Datastore index using `gcloud datastore indexes create`.
- C. Install the `google-cloud-sdk-datastore-emulator` component using the `apt get install` command.
- D. Install the `cloud-datastore-emulator` component using the `gcloud components install` command.

Correct Answer: C

<https://cloud.google.com/sdk/docs/downloads-apt-get>

The question is not about the datastore command itself but from where we should run the update command on the Ubuntu to install the component.