



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<sup>Q&As</sup>

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 development project with appropriate IAM roles defined. You are creating a production project and want to have the same IAM roles on the new project, using the fewest possible steps. What should you do?

- A. Use gcloud iam roles copy and specify the production project as the destination project.
- B. Use gcloud iam roles copy and specify your organization as the destination organization.
- C. In the Google Cloud Platform Console, use the 'create role from role' functionality.
- D. In the Google Cloud Platform Console, use the 'create role' functionality and select all applicable permissions.

Correct Answer: A

Reference: <https://cloud.google.com/sdk/gcloud/reference/iam/roles/copy> It's at project level and iam role copy

---

### QUESTION 2

Your engineers need to pass database credentials to a Kubernetes Pod. The YAML they're using looks similar to the following:

```
apiVersion: "extensions/v1beta1"
```

```
kind: "Deployment"
```

```
metadata:
```

```
name: "products-service"
```

```
namespace: "default"
```

```
labels:
```

```
app: "products-service"
```

```
spec:
```

```
replicas: 3
```

```
selector:
```

```
matchLabels:
```

```
app: "products-service"
```

```
template:
```

```
metadata:
```

```
labels:
```

```
app: "products-service"
```



spec:

containers:

-

name: "products" image: "gcr.io/find-seller-app-dev/products:latest" env:

-

name: "database\_user" value: "admin"

-

name: "database\_password" value: "TheB3stP@ssW0rd"

What is Google's recommended best practice for working with sensitive information inside of Kubernetes?

- A. Store the credentials in a ConfigMap.
- B. Mount the credentials in a volume.
- C. Use an environment variable.
- D. Store the credentials in a Secret.

Correct Answer: D

---

### QUESTION 3

You have deployed an application on a single Compute Engine instance. The application writes logs to disk. Users start reporting errors with the application. You want to diagnose the problem. What should you do?

- A. Navigate to Cloud Logging and view the application logs.
- B. Connect to the instance's serial console and read the application logs.
- C. Configure a Health Check on the instance and set a Low Healthy Threshold value.
- D. Install and configure the Cloud Logging Agent and view the logs from Cloud Logging.

Correct Answer: D

Reference: <https://cloud.google.com/error-reporting/docs/setup/compute-engine>

---

### QUESTION 4

You are performing a monthly security check of your Google Cloud environment and want to know who has access to view data stored in your Google Cloud Project. What should you do?

- A. Enable Audit Logs for all APIs that are related to data storage.
- B. Review the IAM permissions for any role that allows for data access.



C. Review the Identity-Aware Proxy settings for each resource.

D. Create a Data Loss Prevention job.

Correct Answer: B

Reference: <https://cloud.google.com/compute/docs/access>

## QUESTION 5

You have a large 5-TB AVRO file stored in a Cloud Storage bucket. Your analysts are proficient only in SQL and need access to the data stored in this file. You want to find a cost-effective way to complete their request as soon as possible. What should you do?

A. Load data in Cloud Datastore and run a SQL query against it.

B. Create a BigQuery table and load data in BigQuery. Run a SQL query on this table and drop this table after you complete your request.

C. Create external tables in BigQuery that point to Cloud Storage buckets and run a SQL query on these external tables to complete your request.

D. Create a Hadoop cluster and copy the AVRO file to NDfs by compressing it. Load the file in a hive table and provide access to your analysts so that they can run SQL queries.

Correct Answer: C

Breaking down the question into key points

1.

5-TB AVRO file stored in a Cloud Storage bucket.

2.

Analysts are proficient only in SQL

3.

cost-effective way to complete their request as soon as possible

A. ....Load data in Cloud Datastore... (Not Correct because Cloud Datastore is not a good option to run SQL Queries)

B. ...Load data in BigQuery.... (Not Cost Effective because loading the data which is already present in the bucket into BigQuery again is expensive)

C. Create external tables in BigQuery that point to Cloud Storage buckets and run a SQL query on these external tables to complete your request. (This is the right answer as it meets all the requirements from the question)

D. Create a Hadoop cluster and copy the AVRO file to NDfs by compressing it. Load the file in a hive table and provide access to your analysts so that they can run SQL queries. (Too roundabout and indirect. Not the right option)

[ASSOCIATE-CLOUD-](#)

[ASSOCIATE-CLOUD-](#)

[ASSOCIATE-CLOUD-](#)



VCE & PDF

Pass4itSure.com

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

---

[ENGINEER PDF Dumps](#)

[ENGINEER Study Guide](#)

[ENGINEER Braindumps](#)