



VCE & PDF

PassApply.com

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

2024 Latest passapply 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.passapply.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

Your company uses Cloud Storage to store application backup files for disaster recovery purposes. You want to follow Google's recommended practices. Which storage option should you use?

- A. Multi-Regional Storage
- B. Regional Storage
- C. Nearline Storage
- D. Coldline Storage

Correct Answer: D

Reference: <https://cloud.google.com/storage/docs/storage-classes#nearline>

Coldline is a Cloud Storage class designed for long-term archival and disaster recovery. Coldline is perfect for the archival needs of big data or multimedia content, allowing businesses to archive years of data. Coldline provides fast and

instant (millisecond) access to data and changes the way that companies think about storing and accessing their cold data.

Google Cloud Storage Multi-Regional is a highly available and geo-redundant storage class. It's the best storage class for business continuity, or for serving multimedia content to geographically distributed users.

Google Cloud Storage Regional is a highly available storage class redundant within a single region. It's ideal for pairing storage and compute resources within a region, to deliver low end-to-end latency and high throughput for workloads such

as data transcoding or big data analytics workloads

---

### QUESTION 2

Your team uses a third-party monitoring solution. They've asked you to deploy it to the nodes in your Kubernetes Engine Cluster. What's the best way to do that?

- A. Deploy the monitoring pod as a DaemonSet.
- B. Deploy the monitoring pod as a Deployment.
- C. Use Deployment Manager to deploy the monitoring solution.
- D. Connect to each node via SSH and install the monitoring solution.

Correct Answer: A

---

### QUESTION 3

Your team is using Linux instances on Google Cloud. You need to ensure that your team logs in to these instances in



the most secure and cost efficient way. What should you do?

- A. Attach a public IP to the instances and allow incoming connections from the internet on port 22 for SSH.
- B. Use the `gcloud compute ssh` command with the `--tunnel-through-iap` flag. Allow ingress traffic from the IP range 35.235.240.0/20 on port 22.
- C. Use a third party tool to provide remote access to the instances.
- D. Create a bastion host with public internet access. Create the SSH tunnel to the instance through the bastion host.

Correct Answer: B

<https://cloud.google.com/compute/docs/connect/ssh-using-iap#gcloud>

---

#### QUESTION 4

Your team is developing a product catalog that allows end users to search and filter. The full catalog of products consists of about 500 products. The team doesn't have any experience with SQL, or schema migrations, so they're considering a

NoSQL option.

Which database service would work best?

- A. Cloud SQL
- B. Cloud Memorystore
- C. Bigtable
- D. Cloud Datastore

Correct Answer: D

---

#### QUESTION 5

For analysis purposes, you need to send all the logs from all of your Compute Engine instances to a BigQuery dataset called `platform-logs`. You have already installed the Stackdriver Logging agent on all the instances. You want to minimize cost. What should you do?

- A. 1. Give the BigQuery Data Editor role on the `platform-logs` dataset to the service accounts used by your instances.  
2. Update your instances' metadata to add the following value: `logs-destination: bq://platform-logs`.
- B. 1. In Stackdriver Logging, create a logs export with a Cloud Pub/Sub topic called `logs` as a sink.  
2.  
Create a Cloud Function that is triggered by messages in the `logs` topic.
- 3.



Configure that Cloud Function to drop logs that are not from Compute Engine and to insert Compute Engine logs in the platform-logs dataset.

C. 1. In Stackdriver Logging, create a filter to view only Compute Engine logs.

2.

Click Create Export.

3.

Choose BigQuery as Sink Service, and the platform-logs dataset as Sink Destination.

D. 1. Create a Cloud Function that has the BigQuery User role on the platform-logs dataset.

2.

Configure this Cloud Function to create a BigQuery Job that executes this query:

```
INSERT INTO dataset.platform-logs (timestamp, log)
```

```
SELECT timestamp, log FROM compute.logs
```

```
WHERE timestamp > DATE_SUB(CURRENT_DATE(), INTERVAL 1 DAY)
```

3.

Use Cloud Scheduler to trigger this Cloud Function once a day.

Correct Answer: C

C. is correct, Sinks control how Cloud Logging routes logs. Using sinks, you can route some or all of your logs to supported destinations.

[ASSOCIATE-CLOUD-ENGINEER VCE Dumps](#)

[ASSOCIATE-CLOUD-ENGINEER Exam Questions](#)

[ASSOCIATE-CLOUD-ENGINEER Braindumps](#)