

➤ **Vendor: Amazon**➤ **Exam Code: Associate-Cloud-Engineer**➤ **Exam Name: Associate Cloud Engineer**➤ **New Updated Questions from Braindump2go (Updated in August/2020)****Visit Braindump2go and Download Full Version Associate-Cloud-Engineer Exam Dumps****QUESTION 174**

Your Dataproc cluster runs in a single Virtual Private Cloud (VPC) network in a single subnet with range 172.16.20.128/25. There are no private IP addresses available in the VPC network. You want to add new VMs to communicate with your cluster using the minimum number of steps. What should you do?

- A. Modify the existing subnet range to 172.16.20.0/24.
- B. Create a new Secondary IP Range in the VPC and configure the VMs to use that range.
- C. Create a new VPC network for the VMs.
Enable VPC Peering between the VMs' VPC network and the Dataproc cluster VPC network.
- D. Create a new VPC network for the VMs with a subnet of 172.32.0.0/16.
Enable VPC network Peering between the Dataproc VPC network and the VMs VPC network.
Configure a custom Route exchange.

Answer: C**Explanation:**

A subnet has a single primary IP address range and, optionally, one or more secondary IP address ranges. For each subnet IP address range, Google Cloud creates a subnet route. When you use VPC Network Peering, Google Cloud always exchanges the subnet routes that don't use privately reused public IP addresses between the two peered networks. If firewall rules in each network permit communication, VM instances in one network can communicate with instances in the peered network.

Reference: <https://cloud.google.com/vpc/docs/vpc-peering>

QUESTION 175

You manage an App Engine Service that aggregates and visualizes data from BigQuery. The application is deployed with the default App Engine Service account. The data that needs to be visualized resides in a different project managed by another team. You do not have access to this project, but you want your application to be able to read data from the BigQuery dataset. What should you do?

- A. Ask the other team to grant your default App Engine Service account the role of BigQuery Job User.
- B. Ask the other team to grant your default App Engine Service account the role of BigQuery Data Viewer.
- C. In Cloud IAM of your project, ensure that the default App Engine service account has the role of BigQuery Data Viewer.
- D. In Cloud IAM of your project, grant a newly created service account from the other team the role of BigQuery Job User in your project.

Answer: B**Explanation:****[Associate-Cloud-Engineer Exam Dumps](#) [Associate-Cloud-Engineer Exam Questions](#)****[Associate-Cloud-Engineer PDF Dumps](#) [Associate-Cloud-Engineer VCE Dumps](#)****<https://www.braindump2go.com/associate-cloud-engineer.html>**

The Owner, Editor, and Viewer primitive roles include the BigQuery Admin (roles/bigquery.dataOwner), BigQuery Data Editor (roles/bigquery.dataEditor), and BigQuery Data Viewer (roles/bigquery.dataViewer) roles, respectively. This means the Owner, Editor, and Viewer primitive roles have BigQuery access as defined for the respective BigQuery roles.

Reference: <https://cloud.google.com/bigquery/docs/access-control>

QUESTION 176

You need to create a copy of a custom Compute Engine virtual machine (VM) to facilitate an expected increase in application traffic due to a business acquisition. What should you do?

- A. Create a Compute Engine snapshot of your base VM.
Create your images from that snapshot.
- B. Create a Compute Engine snapshot of your base VM.
Create your instances from that snapshot.
- C. Create a custom Compute Engine image from a snapshot.
Create your images from that image.
- D. Create a custom Compute Engine image from a snapshot.
Create your instances from that image.

Answer: D

Explanation:

A custom image belongs only to your project. To create an instance with a custom image, you must first have a custom image.

Reference: <https://cloud.google.com/compute/docs/instances/create-start-instance>

QUESTION 177

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.

Answer: A

Explanation:

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

QUESTION 178

An application generates daily reports in a Compute Engine virtual machine (VM). The VM is in the project corp-iot-insights. Your team operates only in the project corp-aggregate-reports and needs a copy of the daily exports in the bucket corp-aggregate-reports-storage. You want to configure access so that the daily reports from the VM are available in the bucket corp-aggregate-reports-storage and use as few steps as possible while following Google-recommended practices. What should you do?

- A. Move both projects under the same folder.
- B. Grant the VM Service Account the role Storage Object Creator on corp-aggregate-reports-storage.
- C. Create a Shared VPC network between both projects.
Grant the VM Service Account the role Storage Object Creator on corp-iot-insights.
- D. Make corp-aggregate-reports-storage public and create a folder with a pseudo-randomized suffix name.
Share the folder with the IoT team.

[Associate-Cloud-Engine Exam Dumps](#) [Associate-Cloud-Engine Exam Questions](#)

[Associate-Cloud-Engine PDF Dumps](#) [Associate-Cloud-Engine VCE Dumps](#)

<https://www.braindump2go.com/associate-cloud-engineer.html>

Answer: A

Explanation:

<https://cloud.google.com/billing/docs/onboarding-checklist>

QUESTION 179

You built an application on your development laptop that uses Google Cloud services. Your application uses Application Default Credentials for authentication and works fine on your development laptop. You want to migrate this application to a Compute Engine virtual machine (VM) and set up authentication using Google-recommended practices and minimal changes. What should you do?

- A. Assign appropriate access for Google services to the service account used by the Compute Engine VM.
- B. Create a service account with appropriate access for Google services, and configure the application to use this account.
- C. Store credentials for service accounts with appropriate access for Google services in a config file, and deploy this config file with your application.
- D. Store credentials for your user account with appropriate access for Google services in a config file, and deploy this config file with your application.

Answer: B

Explanation: <https://cloud.google.com/compute/docs/access/create-enable-service-accounts-for-instances>

QUESTION 180

You need to create a Compute Engine instance in a new project that doesn't exist yet. What should you do?

- A. Using the Cloud SDK, create a new project, enable the Compute Engine API in that project, and then create the instance specifying your new project.
- B. Enable the Compute Engine API in the Cloud Console, use the Cloud SDK to create the instance, and then use the -pproject flag to specify a new project.
- C. Using the Cloud SDK, create the new instance, and use the -pproject flag to specify the new project.
Answer yes when prompted by Cloud SDK to enable the Compute Engine API.
- D. Enable the Compute Engine API in the Cloud Console. Go to the Compute Engine section of the Console to create a new instance, and look for the Create In A New Project option in the creation form.

Answer: B

QUESTION 181

Your company runs one batch process in an on-premises server that takes around 30 hours to complete. The task runs monthly, can be performed offline, and must be restarted if interrupted. You want to migrate this workload to the cloud while minimizing cost. What should you do?

- A. Migrate the workload to a Compute Engine Preemptible VM.
- B. Migrate the workload to a Google Kubernetes Engine cluster with Preemptible nodes.
- C. Migrate the workload to a Compute Engine VM.
Start and stop the instance as needed.
- D. Create an Instance Template with Preemptible VMs On.
Create a Managed Instance Group from the template and adjust Target CPU Utilization.
Migrate the workload.

Answer: B

Explanation:

<https://cloud.google.com/kubernetes-engine/docs/tutorials/migrating-node-pool>

[Associate-Cloud-Engine Exam Dumps](#) **[Associate-Cloud-Engine Exam Questions](#)**

[Associate-Cloud-Engine PDF Dumps](#) **[Associate-Cloud-Engine VCE Dumps](#)**

<https://www.braindump2go.com/associate-cloud-engineer.html>

QUESTION 182

You are developing a new application and are looking for a Jenkins installation to build and deploy your source code. You want to automate the installation as quickly and easily as possible. What should you do?

- A. Deploy Jenkins through the Google Cloud Marketplace.
- B. Create a new Compute Engine instance.
Run the Jenkins executable.
- C. Create a new Kubernetes Engine cluster.
Create a deployment for the Jenkins image.
- D. Create an instance template with the Jenkins executable.
Create a managed instance group with this template.

Answer: C

Explanation:

<https://cloud.google.com/solutions/jenkins-on-kubernetes-engine>

QUESTION 183

You have downloaded and installed the gcloud command line interface (CLI) and have authenticated with your Google Account. Most of your Compute Engine instances in your project run in the europe-west1-d zone. You want to avoid having to specify this zone with each CLI command when managing these instances. What should you do?

- A. Set the europe-west1-d zone as the default zone using the gcloud config subcommand.
- B. In the Settings page for Compute Engine under Default location, set the zone to europe-west1-d.
- C. In the CLI installation directory, create a file called default.conf containing zone=europe-west1-d.
- D. Create a Metadata entry on the Compute Engine page with key compute/zone and value europe-west1-d

Answer: C

Explanation:

<https://cloud.google.com/compute/docs/gcloud-compute>

QUESTION 184

The core business of your company is to rent out construction equipment at a large scale. All the equipment that is being rented out has been equipped with multiple sensors that send event information every few seconds. These signals can vary from engine status, distance traveled, fuel level, and more. Customers are billed based on the consumption monitored by these sensors. You expect high throughput ?up to thousands of events per hour per device ?and need to retrieve consistent data based on the time of the event. Storing and retrieving individual signals should be atomic. What should you do?

- A. Create a file in Cloud Storage per device and append new data to that file.
- B. Create a file in Cloud Filestore per device and append new data to that file.
- C. Ingest the data into Datastore. Store data in an entity group based on the device.
- D. Ingest the data into Cloud Bigtable. Create a row key based on the event timestamp.

Answer: C

[Associate-Cloud-Engine Exam Dumps](#) **[Associate-Cloud-Engine Exam Questions](#)**

[Associate-Cloud-Engine PDF Dumps](#) **[Associate-Cloud-Engine VCE Dumps](#)**

<https://www.braindump2go.com/associate-cloud-engineer.html>