

➤ **Vendor: Google**

➤ **Exam Code: Associate-Cloud-Engineer**

➤ **Exam Name: Associate Cloud Engineer**

➤ **New Updated Questions from [Braindump2go](#) (Updated in [January/2021](#))**

**[Visit Braindump2go and Download Full Version Associate-Cloud-Engineer Exam Dumps](#)**

#### **QUESTION 228**

You received a JSON file that contained a private key of a Service Account in order to get access to several resources in a Google Cloud project. You downloaded and installed the Cloud SDK and want to use this private key for authentication and authorization when performing gcloud commands. What should you do?

- A. Use the command `gcloud auth login` and point it to the private key
- B. Use the command `gcloud auth activate-service-account` and point it to the private key
- C. Place the private key file in the installation directory of the Cloud SDK and rename it to "credentials ison"
- D. Place the private key file in your home directory and rename it to "GOOGLE\_APPLICATION\_CREDENTIALS".

**Answer: B**

**Explanation:**

<https://cloud.google.com/sdk/docs/authorizing>

#### **QUESTION 229**

Your company has an internal application for managing transactional orders. The application is used exclusively by employees in a single physical location. The application requires strong consistency, fast queries, and ACID guarantees for multi-table transactional updates. The first version of the application is implemented in PostgreSQL, and you want to deploy it to the cloud with minimal code changes. Which database is most appropriate for this application?

- A. BigQuery
- B. Cloud SQL
- C. Cloud Spanner
- D. Cloud Datastore

**Answer: B**

#### **QUESTION 230**

You are using Data Studio to visualize a table from your data warehouse that is built on top of BigQuery. Data is appended to the data warehouse during the day. At night, the daily summary is recalculated by overwriting the table. You just noticed that the charts in Data Studio are broken, and you want to analyze the problem. What should you do?

- A. Use the BigQuery interface to review the nightly Job and look for any errors
- B. Review the Error Reporting page in the Cloud Console to find any errors.
- C. In Cloud Logging create a filter for your Data Studio report
- D. Use Cloud Debugger to find out why the data was not refreshed correctly

**[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>**

**Answer: D**

**QUESTION 231**

You have created an application that is packaged into a Docker image. You want to deploy the Docker image as a workload on Google Kubernetes Engine. What should you do?

- A. Upload the image to Cloud Storage and create a Kubernetes Service referencing the image
- B. Upload the image to Cloud Storage and create a Kubernetes Deployment referencing the image
- C. Upload the image to Container Registry and create a Kubernetes Service referencing the image.
- D. Upload the image to Container Registry and create a Kubernetes Deployment referencing the image

**Answer: C**

**QUESTION 232**

You have just created a new project which will be used to deploy a globally distributed application. You will use Cloud Spanner for data storage. You want to create a Cloud Spanner instance. You want to perform the first step in preparation of creating the instance. What should you do?

- A. Grant yourself the IAM role of Cloud Spanner Admin
- B. Create a new VPC network with subnetworks in all desired regions
- C. Configure your Cloud Spanner instance to be multi-regional
- D. Enable the Cloud Spanner API

**Answer: C**

**QUESTION 233**

You need to manage a Cloud Spanner Instance for best query performance. Your instance in production runs in a single Google Cloud region. You need to improve performance in the shortest amount of time. You want to follow Google best practices for service configuration. What should you do?

- A. Create an alert in Cloud Monitoring to alert when the percentage of high priority CPU utilization reaches 45%.  
If you exceed this threshold, add nodes to your instance.
- B. Create an alert in Cloud Monitoring to alert when the percentage of high priority CPU utilization reaches 45%.  
Use database query statistics to identify queries that result in high CPU usage, and then rewrite those queries to optimize their resource usage
- C. Create an alert in Cloud Monitoring to alert when the percentage of high priority CPU utilization reaches 65%.  
If you exceed this threshold, add nodes to your instance
- D. Create an alert in Cloud Monitoring to alert when the percentage of high priority CPU utilization reaches 65%.  
Use database query statistics to identify queries that result in high CPU usage, and then rewrite those queries to optimize their resource usage.

**Answer: A**

**QUESTION 234**

You are working with a Cloud SQL MySQL database at your company. You need to retain a month-end copy of the database for three years for audit purposes. What should you do?

- A. Save file automatic first-of-the- month backup for three years Store the backup file in an Archive

**[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>**

class Cloud Storage bucket

- B. Convert the automatic first-of-the-month backup to an export file Write the export file to a Coldline class Cloud Storage bucket
- C. Set up an export job for the first of the month Write the export file to an Archive class Cloud Storage bucket
- D. Set up an on-demand backup for the first of the month Write the backup to an Archive class Cloud Storage bucket

**Answer: A**

#### **QUESTION 235**

You are working for a hospital that stores its medical images in an on-premises data room. The hospital wants to use Cloud Storage for archival storage of these images. The hospital wants an automated process to upload any new medical images to Cloud Storage. You need to design and implement a solution. What should you do?

- A. Deploy a Dataflow job from the batch template "Datastore to Cloud Storage" Schedule the batch job on the desired interval
- B. In the Cloud Console, go to Cloud Storage Upload the relevant images to the appropriate bucket
- C. Create a script that uses the gsutil command line interface to synchronize the on-premises storage with Cloud Storage Schedule the script as a cron job
- D. Create a Pub/Sub topic, and enable a Cloud Storage trigger for the Pub/Sub topic. Create an application that sends all medical images to the Pub/Sub topic

**Answer: A**

#### **QUESTION 236**

You have been asked to create robust Virtual Private Network (VPN) connectivity between a new Virtual Private Cloud (VPC) and a remote site. Key requirements include dynamic routing, a shared address space of 10.19.0.1/22, and no overprovisioning of tunnels during a failover event. You want to follow Google-recommended practices to set up a high availability Cloud VPN. What should you do?

- A. Use a custom mode VPC network, configure static routes, and use active/passive routing
- B. Use an automatic mode VPC network, configure static routes, and use active/active routing
- C. Use a custom mode VPC network use Cloud Router border gateway protocol (BGP) routes, and use active/passive routing
- D. Use an automatic mode VPC network, use Cloud Router border gateway protocol (BGP) routes and configure policy-based routing

**Answer: D**

#### **QUESTION 237**

You are creating an application that will run on Google Kubernetes Engine. You have identified MongoDB as the most suitable database system for your application and want to deploy a managed MongoDB environment that provides a support SLA. What should you do?

- A. Create a Cloud Bigtable cluster and use the HBase API
- B. Deploy MongoDB Alias from the Google Cloud Marketplace
- C. Download a MongoDB installation package and run it on Compute Engine instances
- D. Download a MongoDB installation package, and run it on a Managed Instance Group

**Answer: D**

#### **QUESTION 238**

**[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>**

You need to configure optimal data storage for files stored in Cloud Storage for minimal cost. The files are used in a mission-critical analytics pipeline that is used continually. The users are in Boston, MA (United States). What should you do?

- A. Configure regional storage for the region closest to the users.  
Configure a Nearline storage class
- B. Configure regional storage for the region closest to the users.  
Configure a Standard storage class
- C. Configure dual-regional storage for the dual region closest to the users.  
Configure a Nearline storage class
- D. Configure dual-regional storage for the dual region closest to the users.  
Configure a Standard storage class

**Answer: B**

#### **QUESTION 239**

You are developing a new web application that will be deployed on Google Cloud Platform. As part of your release cycle, you want to test updates to your application on a small portion of real user traffic. The majority of the users should still be directed towards a stable version of your application. What should you do?

- A. Deploy the application on App Engine.  
For each update, create a new version of the same service.  
Configure traffic splitting to send a small percentage of traffic to the new version
- B. Deploy the application on App Engine.  
For each update, create a new service.  
Configure traffic splitting to send a small percentage of traffic to the new service.
- C. Deploy the application on Kubernetes Engine.  
For a new release, update the deployment to use the new version
- D. Deploy the application on Kubernetes Engine.  
For a new release, create a new deployment for the new version. Update the service to use the new deployment.

**Answer: A**

#### **QUESTION 240**

Your company developed a mobile game that is deployed on Google Cloud. Gamers are connecting to the game with their personal phones over the Internet. The game sends UDP packets to update the servers about the gamers' actions while they are playing in multiplayer mode. Your game backend can scale over multiple virtual machines (VMs), and you want to expose the VMs over a single IP address. What should you do?

- A. Configure an SSL Proxy load balancer in front of the application servers.
- B. Configure an Internal UDP load balancer in front of the application servers.
- C. Configure an External HTTP(s) load balancer in front of the application servers.
- D. Configure an External Network load balancer in front of the application servers.

**Answer: A**

**Explanation:**

<https://cloud.google.com/solutions/connecting-securely>

#### **QUESTION 241**

You are working with a Cloud SQL MySQL database at your company. You need to retain a month-end copy of the database for three years for audit purposes. What should you do?

- A. Set up an export job for the first of the month.

**[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>**

Write the export file to an Archive class Cloud Storage bucket.

- B. Save the automatic first-of-the-month backup for three years.  
Store the backup file in an Archive class Cloud Storage bucket.
- C. Set up an on-demand backup for the first of the month.  
Write the backup to an Archive class Cloud Storage bucket.
- D. Convert the automatic first-of-the-month backup to an export file.  
Write the export file to a Coldline class Cloud Storage bucket.

**Answer: B**

**QUESTION 242**

You are monitoring an application and receive user feedback that a specific error is spiking. You notice that the error is caused by a Service Account having insufficient permissions. You are able to solve the problem but want to be notified if the problem recurs. What should you do?

- A. In the Log Viewer, filter the logs on severity 'Error' and the name of the Service Account.
- B. Create a sink to BigQuery to export all the logs.  
Create a Data Studio dashboard on the exported logs.
- C. Create a custom log-based metrics for the specific error to be used in an Alerting Policy.
- D. Grant Project Owner access to the Service Account.

**Answer: A**

**Explanation:**

<https://cloud.google.com/logging/docs/view/advanced-queries8E6BA412E7DB6A14A62CC68E5EB6DAE3>

**[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>**