

➤ **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 185**

You are asked to set up application performance monitoring on Google Cloud projects A, B, and C as a single pane of glass. You want to monitor CPU, memory, and disk. What should you do?

- A. Enable API and then share charts from project A, B, and C.
- B. Enable API and then give the metrics.reader role to projects A, B, and C.
- C. Enable API and then use default dashboards to view all projects in sequence.
- D. Enable API, create a workspace under project A, and then add project B and C.

**Answer: C****QUESTION 186**

You created several resources in multiple Google Cloud projects. All projects are linked to different billing accounts. To better estimate future charges, you want to have a single visual representation of all costs incurred. You want to include new cost data as soon as possible. What should you do?

- A. Configure Billing Data Export to BigQuery and visualize the data in Data Studio.
- B. Visit the Cost Table page to get a CSV export and visualize it using Data Studio.
- C. Fill all resources in the Pricing Calculator to get an estimate of the monthly cost.
- D. Use the Reports view in the Cloud Billing Console to view the desired cost information.

**Answer: A****Explanation:**<https://cloud.google.com/billing/docs/how-to/visualize-data>**QUESTION 187**

Your company has workloads running on Compute Engine and on-premises. The Google Cloud Virtual Private Cloud (VPC) is connected to your WAN over a Virtual Private Network (VPN). You need to deploy a new Compute Engine instance and ensure that no public Internet traffic can be routed to it. What should you do?

- A. Create the instance without a public IP address.
- B. Create the instance with Private Google Access enabled.
- C. Create a deny-all egress firewall rule on the VPC network.
- D. Create a route on the VPC to route all traffic to the instance over the VPN tunnel.

**Answer: B****Explanation:**

Get private access to Google services, such as storage, big data, analytics, or machine learning, without having to give your service a public IP address.

**[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>**

Reference: <https://cloud.google.com/vpc>

**QUESTION 188**

Your team maintains the infrastructure for your organization. The current infrastructure requires changes. You need to share your proposed changes with the rest of the team. You want to follow Google's recommended best practices. What should you do?

- A. Use Deployment Manager templates to describe the proposed changes and store them in a Cloud Storage bucket.
- B. Use Deployment Manager templates to describe the proposed changes and store them in Cloud Source Repositories.
- C. Apply the change in a development environment, run `gcloud compute instances list`, and then save the output in a shared Storage bucket.
- D. Apply the change in a development environment, run `gcloud compute instances list`, and then save the output in Cloud Source Repositories.

**Answer: B**

**QUESTION 189**

You have a Compute Engine instance hosting an application used between 9 AM and 6 PM on weekdays. You want to back up this instance daily for disaster recovery purposes. You want to keep the backups for 30 days. You want the Google-recommended solution with the least management overhead and the least number of services. What should you do?

- A.
  - 1. Update your instances' metadata to add the following value: `snapshot-schedule: 0 1 * * *`
  - 2. Update your instances' metadata to add the following value: `snapshot-retention: 30`
- B.
  - 1. In the Cloud Console, go to the Compute Engine Disks page and select your instance's disk.
  - 2. In the Snapshot Schedule section, select Create Schedule and configure the following parameters:
    - Schedule frequency: Daily
    - Start time: 1:00 AM ?2:00 AM
    - Autodelete snapshots after 30 days
- C.
  - 1. Create a Cloud Function that creates a snapshot of your instance's disk.
  - 2. Create a Cloud Function that deletes snapshots that are older than 30 days.
  - 3. Use Cloud Scheduler to trigger both Cloud Functions daily at 1:00 AM.
- D.
  - 1. Create a bash script in the instance that copies the content of the disk to Cloud Storage.
  - 2. Create a bash script in the instance that deletes data older than 30 days in the backup Cloud Storage bucket.
  - 3. Configure the instance's crontab to execute these scripts daily at 1:00 AM.

**Answer: B**

**QUESTION 190**

Your existing application running in Google Kubernetes Engine (GKE) consists of multiple pods running on four GKE `n1-standard-2` nodes. You need to deploy additional pods requiring `n2-highmem-16` nodes without any downtime. What should you do?

- A. Use `gcloud container clusters upgrade`. Deploy the new services.
- B. Create a new Node Pool and specify machine type `n2-highmem-16`. Deploy the new pods.
- C. Create a new cluster with `n2-highmem-16` nodes. Redeploy the pods and delete the old cluster.
- D. Create a new cluster with both `n1-standard-2` and `n2-highmem-16` nodes.

**[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>**

Redeploy the pods and delete the old cluster.

**Answer: B**

**QUESTION 191**

You have an application that uses Cloud Spanner as a database backend to keep current state information about users. Cloud Bigtable logs all events triggered by users. You export Cloud Spanner data to Cloud Storage during daily backups. One of your analysts asks you to join data from Cloud Spanner and Cloud Bigtable for specific users. You want to complete this ad hoc request as efficiently as possible. What should you do?

- A. Create a dataflow job that copies data from Cloud Bigtable and Cloud Storage for specific users.
- B. Create a dataflow job that copies data from Cloud Bigtable and Cloud Spanner for specific users.
- C. Create a Cloud Dataproc cluster that runs a Spark job to extract data from Cloud Bigtable and Cloud Storage for specific users.
- D. Create two separate BigQuery external tables on Cloud Storage and Cloud Bigtable. Use the BigQuery console to join these tables through user fields, and apply appropriate filters.

**Answer: B**

**QUESTION 192**

You are hosting an application from Compute Engine virtual machines (VMs) in us-central1-a. You want to adjust your design to support the failure of a single Compute Engine zone, eliminate downtime, and minimize cost. What should you do?

- A. – Create Compute Engine resources in us-central1-b.
  - Balance the load across both us-central1-a and us-central1-b.
- B. – Create a Managed Instance Group and specify us-central1-a as the zone.
  - Configure the Health Check with a short Health Interval.
- C. – Create an HTTP(S) Load Balancer.
  - Create one or more global forwarding rules to direct traffic to your VMs.
- D. – Perform regular backups of your application.
  - Create a Cloud Monitoring Alert and be notified if your application becomes unavailable.
  - Restore from backups when notified.

**Answer: B**

**Explanation:**

<https://github.com/GoogleCloudPlatform/puppet-google-compute>

**QUESTION 193**

A colleague handed over a Google Cloud Platform project for you to maintain. As part of a security checkup, you want to review who has been granted the Project Owner role. What should you do?

- A. In the console, validate which SSH keys have been stored as project-wide keys.
- B. Navigate to Identity-Aware Proxy and check the permissions for these resources.
- C. Enable Audit Logs on the IAM & admin page for all resources, and validate the results.
- D. Use the command `gcloud projects get-iam-policy` to view the current role assignments.

**Answer: A**

**Explanation:**

<https://cloud.google.com/compute/docs/instances/adding-removing-ssh-keys>

**QUESTION 194**

You are running multiple VPC-native Google Kubernetes Engine clusters in the same subnet. The IPs available for the nodes are exhausted, and you want to ensure that the clusters can grow in nodes when needed. What should you do?

**[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>**

- A. Create a new subnet in the same region as the subnet being used.
- B. Add an alias IP range to the subnet used by the GKE clusters.
- C. Create a new VPC, and set up VPC peering with the existing VPC.
- D. Expand the CIDR range of the relevant subnet for the cluster.

**Answer: C**

**Explanation:**

To create a VPC peering connection, first create a request to peer with another VPC.

Reference: <https://docs.aws.amazon.com/vpc/latest/peering/vpc-pg.pdf>

#### **QUESTION 195**

You have a batch workload that runs every night and uses a large number of virtual machines (VMs). It is fault-tolerant and can tolerate some of the VMs being terminated. The current cost of VMs is too high.

What should you do?

- A. Run a test using simulated maintenance events.  
If the test is successful, use preemptible N1 Standard VMs when running future jobs.
- B. Run a test using simulated maintenance events.  
If the test is successful, use N1 Standard VMs when running future jobs.
- C. Run a test using a managed instance group.  
If the test is successful, use N1 Standard VMs in the managed instance group when running future jobs.
- D. Run a test using N1 standard VMs instead of N2.  
If the test is successful, use N1 Standard VMs when running future jobs.

**Answer: B**

**Explanation:**

<https://cloud.google.com/compute/vm-instance-pricing>

#### **QUESTION 196**

You are working with a user to set up an application in a new VPC behind a firewall. The user is concerned about data egress. You want to configure the fewest open egress ports. What should you do?

- A. Set up a low-priority (65534) rule that blocks all egress and a high-priority rule (1000) that allows only the appropriate ports.
- B. Set up a high-priority (1000) rule that pairs both ingress and egress ports.
- C. Set up a high-priority (1000) rule that blocks all egress and a low-priority (65534) rule that allows only the appropriate ports.
- D. Set up a high-priority (1000) rule to allow the appropriate ports.

**Answer: C**

#### **QUESTION 197**

Your company runs its Linux workloads on Compute Engine instances. Your company will be working with a new operations partner that does not use Google Accounts. You need to grant access to the instances to your operations partner so they can maintain the installed tooling. What should you do?

- A. Enable Cloud IAP for the Compute Engine instances, and add the operations partner as a Cloud IAP Tunnel User.
- B. Tag all the instances with the same network tag.  
Create a firewall rule in the VPC to grant TCP access on port 22 for traffic from the operations partner to instances with the network tag.
- C. Set up Cloud VPN between your Google Cloud VPC and the internal network of the operations

**[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>**

partner.

- D. Ask the operations partner to generate SSH key pairs, and add the public keys to the VM instances.

**Answer: B**

**Explanation:**

<https://cloud.google.com/vpc/docs/firewalls>

#### **QUESTION 198**

You have created a code snippet that should be triggered whenever a new file is uploaded to a Cloud Storage bucket. You want to deploy this code snippet. What should you do?

- A. Use App Engine and configure Cloud Scheduler to trigger the application using Pub/Sub.
- B. Use Cloud Functions and configure the bucket as a trigger resource.
- C. Use Google Kubernetes Engine and configure a CronJob to trigger the application using Pub/Sub.
- D. Use Dataflow as a batch job, and configure the bucket as a data source.

**Answer: A**

**Explanation:**

<https://cloud.google.com/scheduler/docs/tut-pub-sub>

#### **QUESTION 199**

You have been asked to set up Object Lifecycle Management for objects stored in storage buckets. The objects are written once and accessed frequently for 30 days. After 30 days, the objects are not read again unless there is a special need. The object should be kept for three years, and you need to minimize cost. What should you do?

- A. Set up a policy that uses Nearline storage for 30 days and then moves to Archive storage for three years.
- B. Set up a policy that uses Standard storage for 30 days and then moves to Archive storage for three years.
- C. Set up a policy that uses Nearline storage for 30 days, then moves the Coldline for one year, and then moves to Archive storage for two years.
- D. Set up a policy that uses Standard storage for 30 days, then moves to Coldline for one year, and then moves to Archive storage for two years.

**Answer: A**

#### **QUESTION 200**

You are storing sensitive information in a Cloud Storage bucket. For legal reasons, you need to be able to record all requests that read any of the stored data. You want to make sure you comply with these requirements. What should you do?

- A. Enable the Identity Aware Proxy API on the project.
- B. Scan the bucket using the Data Loss Prevention API.
- C. Allow only a single Service Account access to read the data.
- D. Enable Data Access audit logs for the Cloud Storage API.

**Answer: D**

**Explanation:**

<https://cloud.google.com/storage/docs/audit-logs>

#### **QUESTION 201**

**[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 are the team lead of a group of 10 developers. You provided each developer with an individual Google Cloud Project that they can use as their personal sandbox to experiment with different Google Cloud solutions. You want to be notified if any of the developers are spending above \$500 per month on their sandbox environment. What should you do?

- A. Create a single budget for all projects and configure budget alerts on this budget.
- B. Create a separate billing account per sandbox project and enable BigQuery billing exports.  
Create a Data Studio dashboard to plot the spending per billing account.
- C. Create a budget per project and configure budget alerts on all of these budgets.
- D. Create a single billing account for all sandbox projects and enable BigQuery billing exports.  
Create a Data Studio dashboard to plot the spending per project.

**Answer: C**

**Explanation:**

<https://cloud.google.com/billing/docs/how-to/budgets>

#### **QUESTION 202**

You are deploying a production application on Compute Engine. You want to prevent anyone from accidentally destroying the instance by clicking the wrong button. What should you do?

- A. Disable the flag "Delete boot disk when instance is deleted."
- B. Enable delete protection on the instance.
- C. Disable Automatic restart on the instance.
- D. Enable Preemptibility on the instance.

**Answer: A**

**Explanation:**

<https://googlecloudplatform.uservoice.com/forums/302595-compute-engine/suggestions/14227521-set-delete-boot-disk-when-instance-is-deleted-to>

#### **QUESTION 203**

Your company uses a large number of Google Cloud services centralized in a single project. All teams have specific projects for testing and development. The DevOps team needs access to all of the production services in order to perform their job. You want to prevent Google Cloud product changes from broadening their permissions in the future. You want to follow Google-recommended practices. What should you do?

- A. Grant all members of the DevOps team the role of Project Editor on the organization level.
- B. Grant all members of the DevOps team the role of Project Editor on the production project.
- C. Create a custom role that combines the required permissions.  
Grant the DevOps team the custom role on the production project.
- D. Create a custom role that combines the required permissions.  
Grant the DevOps team the custom role on the organization level.

**Answer: A**

#### **QUESTION 204**

You are building an application that processes data files uploaded from thousands of suppliers. Your primary goals for the application are data security and the expiration of aged data. You need to design the application to:

- Restrict access so that suppliers can access only their own data.
- Give suppliers write access to data only for 30 minutes.
- Delete data that is over 45 days old.

You have a very short development cycle, and you need to make sure that the application requires minimal maintenance. Which two strategies should you use? (Choose two.)

**[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>**



- A. Build a lifecycle policy to delete Cloud Storage objects after 45 days.
- B. Use signed URLs to allow suppliers limited time access to store their objects.
- C. Set up an SFTP server for your application, and create a separate user for each supplier.
- D. Build a Cloud function that triggers a timer of 45 days to delete objects that have expired.
- E. Develop a script that loops through all Cloud Storage buckets and deletes any buckets that are older than 45 days.

**Answer:** AE

**[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>**