

[August-2020Braindump2go Associate-Cloud-Engineer Exam Dumps(PDF and VCE) Associate-Cloud-Engineer 204Q for 100% Passing Associate-Cloud-Engineer Exam[Q174-Q189]

2020/August Latest Braindump2go Associate-Cloud-Engineer Exam Dumps with PDF and VCE Free Updated Today! Following are some new Associate-Cloud-Engineer Real Exam Questions!QUESTION 174Your 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 175You 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: 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 176You 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 177You 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 178An 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.
Answer: A
Explanation: <https://cloud.google.com/billing/docs/onboarding-checklist>QUESTION 179You 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>

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 `eu-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 `eu-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 `eu-west1-d`.
C. In the CLI installation directory, create a file called `default.conf` containing `zone=eu-west1-d`.
D. Create a Metadata entry on the Compute Engine page with key `compute/zone` and value `eu-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
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.
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

Resources From: 1. 2020 Latest Braindump2go Associate-Cloud-Engineer Exam Dumps (PDF & VCE) Free Share:
<https://www.braindump2go.com/associate-cloud-engineer.html>
2. 2020 Latest Braindump2go Associate-Cloud-Engineer PDF and VCE Dumps Free Share:
https://drive.google.com/drive/folders/1Z0bKeusnOFNlsc3XP3vHM_m_DrdKaFmA?usp=sharing
3. 2020 Free Braindump2go Associate-Cloud-Engineer PDF Download:
[https://www.braindump2go.com/free-online-pdf/Associate-Cloud-Engineer-PDF-Dumps\(185-204\).pdf](https://www.braindump2go.com/free-online-pdf/Associate-Cloud-Engineer-PDF-Dumps(185-204).pdf)
[https://www.braindump2go.com/free-online-pdf/Associate-Cloud-Engineer-VCE-Dumps\(174-184\).pdf](https://www.braindump2go.com/free-online-pdf/Associate-Cloud-Engineer-VCE-Dumps(174-184).pdf)
Free Resources from Braindump2go, We Devoted to Helping You 100% Pass All Exams!