

## 100% Reliable Associate-Cloud-Engineer Exam Dumps Test Pdf Exam Material [Q115-Q130]



### 100% Reliable Microsoft Associate-Cloud-Engineer Exam Dumps Test Pdf Exam Material Based on Official Syllabus Topics of Actual Google Associate-Cloud-Engineer Exam

#### Ensuring Successful Functioning of Cloud Solutions

The aim of this section is to evaluate the expertise of the specialists in managing Compute Engine resources. Furthermore, the applicants are required to possess the skills in managing networking resources, Google Kubernetes Engine resources, Cloud Run & App Engine resources, storage & database solutions. They should also have a good comprehension of how to perform monitoring & logging.

#### NEW QUESTION 115

You are building an application that will run in your data center. The application will use Google Cloud Platform (GCP) services like AutoML. You created a service account that has appropriate access to AutoML. You need to enable authentication to the APIs from your on-premises environment. What should you do?

- \* Use service account credentials in your on-premises application.
- \* Use gcloud to create a key file for the service account that has appropriate permissions.
- \* Set up direct interconnect between your data center and Google Cloud Platform to enable authentication for your on-premises applications.

\* Go to the IAM & admin console, grant a user account permissions similar to the service account permissions, and use this user account for authentication from your data center.

Reference:

<https://cloud.google.com/vision/automl/docs/before-you-begin>

### NEW QUESTION 116

You want to configure autohealing for network load balancing for a group of Compute Engine instances that run in multiple zones, using the fewest possible steps. You need to configure re-creation of VMs if they are unresponsive after 3 attempts of 10 seconds each. What should you do?

\* Create an HTTP load balancer with a backend configuration that references an existing instance group.

Set the health check to healthy (HTTP)

\* Create an HTTP load balancer with a backend configuration that references an existing instance group.

Define a balancing mode and set the maximum RPS to 10.

\* Create a managed instance group. Set the Autohealing health check to healthy (HTTP)

\* Create a managed instance group. Verify that the autoscaling setting is on.

Use separate health checks for load balancing and for autohealing. Health checks for load balancing detect unresponsive instances and direct traffic away from them. Health checks for autohealing detect and recreate failed instances, so they should be less aggressive than load balancing health checks. Using the same health check for these services would remove the distinction between unresponsive instances and failed instances, causing unnecessary latency and unavailability for your users.

<https://cloud.google.com/compute/docs/tutorials/high-availability-autohealing>

### NEW QUESTION 117

You need to select and configure compute resources for a set of batch processing jobs. These jobs take around 2 hours to complete and are run nightly. You want to minimize service costs.

What should you do?

\* Select Google Kubernetes Engine. Use a single-node cluster with a small instance type.

\* Select Google Kubernetes Engine. Use a three-node cluster with micro instance types.

\* Select Compute Engine. Use preemptible VM instances of the appropriate standard machine type.

\* Select Compute Engine. Use VM instance types that support micro bursting.

### NEW QUESTION 118

Your auditor wants to view your organization's use of data in Google Cloud. The auditor is most interested in auditing who accessed data in Cloud Storage buckets. You need to help the auditor access the data they need. What should you do?

\* Turn on Data Access Logs for the buckets they want to audit, and then build a query in the log viewer that filters on Cloud Storage.

\* Assign the appropriate permissions, and then create a Data Studio report on Admin Activity Audit Logs.

\* Assign the appropriate permissions, and the use Cloud Monitoring to review metrics.

\* Use the export logs API to provide the Admin Activity Audit Logs in the format they want.

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

### NEW QUESTION 119

You are using Google Kubernetes Engine with autoscaling enabled to host a new application. You want to expose this new application to the public, using HTTPS on a public IP address. What should you do?

- \* Create a Kubernetes Service of type NodePort for your application, and a Kubernetes Ingress to expose this Service via a Cloud Load Balancer.
  - \* Create a Kubernetes Service of type ClusterIP for your application. Configure the public DNS name of your application using the IP of this Service.
  - \* Create a Kubernetes Service of type NodePort to expose the application on port 443 of each node of the Kubernetes cluster. Configure the public DNS name of your application with the IP of every node of the cluster to achieve load-balancing.
  - \* Create a HAProxy pod in the cluster to load-balance the traffic to all the pods of the application. Forward the public traffic to HAProxy with an iptable rule. Configure the DNS name of your application using the public IP of the node HAProxy is running on.
- Reference:

<https://cloud.google.com/kubernetes-engine/docs/tutorials/http-balancer>

### NEW QUESTION 120

Your development team needs a new Jenkins server for their project. You need to deploy the server using the fewest steps possible. What should you do?

- \* Download and deploy the Jenkins Java WAR to App Engine Standard.
- \* Create a new Compute Engine instance and install Jenkins through the command line interface.
- \* Create a Kubernetes cluster on Compute Engine and create a deployment with the Jenkins Docker image.
- \* Use GCP Marketplace to launch the Jenkins solution.

<https://cloud.google.com/solutions/using-jenkins-for-distributed-builds-on-compute-engine>

### NEW QUESTION 121

You are building an application that will run in your data center. The application will use Google Cloud Platform (GCP) services like AutoML. You created a service account that has appropriate access to AutoML.

You need to enable authentication to the APIs from your on-premises environment. What should you do?

- \* Use service account credentials in your on-premises application.
- \* Use gcloud to create a key file for the service account that has appropriate permissions.
- \* Set up direct interconnect between your data center and Google Cloud Platform to enable authentication for your on-premises applications.
- \* Go to the IAM & admin console, grant a user account permissions similar to the service account permissions, and use this user account for authentication from your data center.

### NEW QUESTION 122

You are building an archival solution for your data warehouse and have selected Cloud Storage to archive your data. Your users need to be able to access this archived data once a quarter for some regulatory requirements. You want to select a cost-efficient option. Which storage option should you use?

- \* Cold Storage
- \* Nearline Storage
- \* Regional Storage
- \* Multi-Regional Storage

Nearline, Coldline, and Archive offer ultra low-cost, highly-durable, highly available archival storage. For data accessed less than once a year, Archive is a cost-effective storage option for long-term preservation of data.

Coldline is also ideal for cold storage-data your business expects to touch less than once a quarter. For warmer storage, choose

Nearline: data you expect to access less than once a month, but possibly multiple times throughout the year. All storage classes are available across all GCP regions and provide unparalleled sub-second access speeds with a consistent API.

### NEW QUESTION 123

You need to set up permissions for a set of Compute Engine instances to enable them to write data into a particular Cloud Storage bucket. You want to follow Google-recommended practices. What should you do?

- \* Create a service account with an access scope. Use the access scope `https://www.googleapis.com/auth/devstorage.write_only`.
- \* Create a service account with an access scope. Use the access scope `https://www.googleapis.com/auth/cloud-platform`.
- \* Create a service account and add it to the IAM role `storage.objectCreator` for that bucket.
- \* Create a service account and add it to the IAM role `storage.objectAdmin` for that bucket.

### NEW QUESTION 124

You have a website hosted on App Engine standard environment. You want 1% of your users to see a new test version of the website. You want to minimize complexity. What should you do?

- \* Deploy the new version in the same application and use the `migrateoption`.
- \* Deploy the new version in the same application and use the `splitsoption` to give a weight of 99 to the current version and a weight of 1 to the new version.
- \* Create a new App Engine application in the same project. Deploy the new version in that application. Use the App Engine library to proxy 1% of the requests to the new version.
- \* Create a new App Engine application in the same project. Deploy the new version in that application.

Configure your network load balancer to send 1% of the traffic to that new application.

### NEW QUESTION 125

Your company is moving from an on-premises environment to Google Cloud Platform (GCP). You have multiple development teams that use Cassandra environments as backend databases. They all need a development environment that is isolated from other Cassandra instances. You want to move to GCP quickly and with minimal support effort. What should you do?

- \* 1. Build an instruction guide to install Cassandra on GCP.
- 2. Make the instruction guide accessible to your developers.
  - \* 1. Advise your developers to go to Cloud Marketplace.
- 2. Ask the developers to launch a Cassandra image for their development work.
  - \* 1. Build a Cassandra Compute Engine instance and take a snapshot of it.
- 2. Use the snapshot to create instances for your developers.
  - \* 1. Build a Cassandra Compute Engine instance and take a snapshot of it.
- 2. Upload the snapshot to Cloud Storage and make it accessible to your developers.
- 3. Build instructions to create a Compute Engine instance from the snapshot so that developers can do it themselves.

### NEW QUESTION 126

You need to update a deployment in Deployment Manager without any resource downtime in the deployment. Which command

should you use?

- \* `gcloud deployment-manager deployments create --config <deployment-config-path>`
  - \* `gcloud deployment-manager deployments update --config <deployment-config-path>`
  - \* `gcloud deployment-manager resources create --config <deployment-config-path>`
  - \* `gcloud deployment-manager resources update --config <deployment-config-path>`
- <https://cloud.google.com/sdk/gcloud/reference/deployment-manager/deployments/update>

### NEW QUESTION 127

You need to create a new billing account and then link it with an existing Google Cloud Platform project.

What should you do?

- \* Verify that you are Project Billing Manager for the GCP project.

Update the existing project to link it to the existing billing account.

- \* Verify that you are Project Billing Manager for the GCP project.

Create a new billing account and link the new billing account to the existing project.

- \* Verify that you are Billing Administrator for the billing account.

Create a new project and link the new project to the existing billing account.

- \* Verify that you are Billing Administrator for the billing account.

Update the existing project to link it to the existing billing account.

### NEW QUESTION 128

A team of data scientists infrequently needs to use a Google Kubernetes Engine (GKE) cluster that you manage. They require GPUs for some long-running, non-restartable jobs. You want to minimize cost. What should you do?

- \* Enable node auto-provisioning on the GKE cluster.
- \* Create a VerticalPodAutoscaler for those workloads.
- \* Create a node pool with preemptible VMs and GPUs attached to those VMs.
- \* Create a node pool of instances with GPUs, and enable autoscaling on this node pool with a minimum size of 1.

### NEW QUESTION 129

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?

- \* Set up a policy that uses Nearline storage for 30 days and then moves to Archive storage for three years.
- \* Set up a policy that uses Standard storage for 30 days and then moves to Archive storage for three years.
- \* 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.
- \* 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.

### NEW QUESTION 130

You are managing several Google Cloud Platform (GCP) projects and need access to all logs for the past 60 days. You want to be able to explore and quickly analyze the log contents. You want to follow Google-recommended practices to obtain the combined

logs for all projects. What should you do?

- \* Navigate to Stackdriver Logging and select resource.labels.project\_id=&#8221;\*&#8221;
- \* Create a Stackdriver Logging Export with a Sink destination to a BigQuery dataset.

Configure the table expiration to 60 days.

- \* Create a Stackdriver Logging Export with a Sink destination to Cloud Storage.

Create a lifecycle rule to delete objects after 60 days.

- \* Configure a Cloud Scheduler job to read from Stackdriver and store the logs in BigQuery.

Configure the table expiration to 60 days.

<https://cloud.google.com/blog/products/gcp/best-practices-for-working-with-google-cloud-audit-logging>

**Free Associate-Cloud-Engineer Dumps are Available for Instant Access:**

<https://www.validexam.com/Associate-Cloud-Engineer-latest-dumps.html>