# Latest Associate-Cloud-Engineer Study Guides 2022 - With Test Engine PDF [Q51-Q65

Latest Associate-Cloud-Engineer Study Guides 2022 - With Test Engine PDF
Get New Associate-Cloud-Engineer Practice Test Questions Answers

Below is the Associate Cloud Engineer Exam Format

Format: Multiple choices, multiple answers

 - Length of Examination: 2 Hours- Passing score: 80%- Language: English, Japanese, Spanish, Portuguese, French, German, and Indonesian.- Number of Questions: 50 **QUESTION 51**

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?

* Use the command gcloud auth loginand point it to the private key.
* Use the command gcloud auth activate-service-accountand point it to the private key.
* Place the private key file in the installation directory of the Cloud SDK and rename it to “credentials.json”.
* Place the private key file in your home directory and rename it to

“GOOGLE_APPLICATION_CREDENTIALS”.
Explanation/Reference: https://cloud.google.com/sdk/docs/authorizing

QUESTION 52

You are building a pipeline to process time-series data. Which Google Cloud Platform services should you put in boxes 1,2,3, and 4?
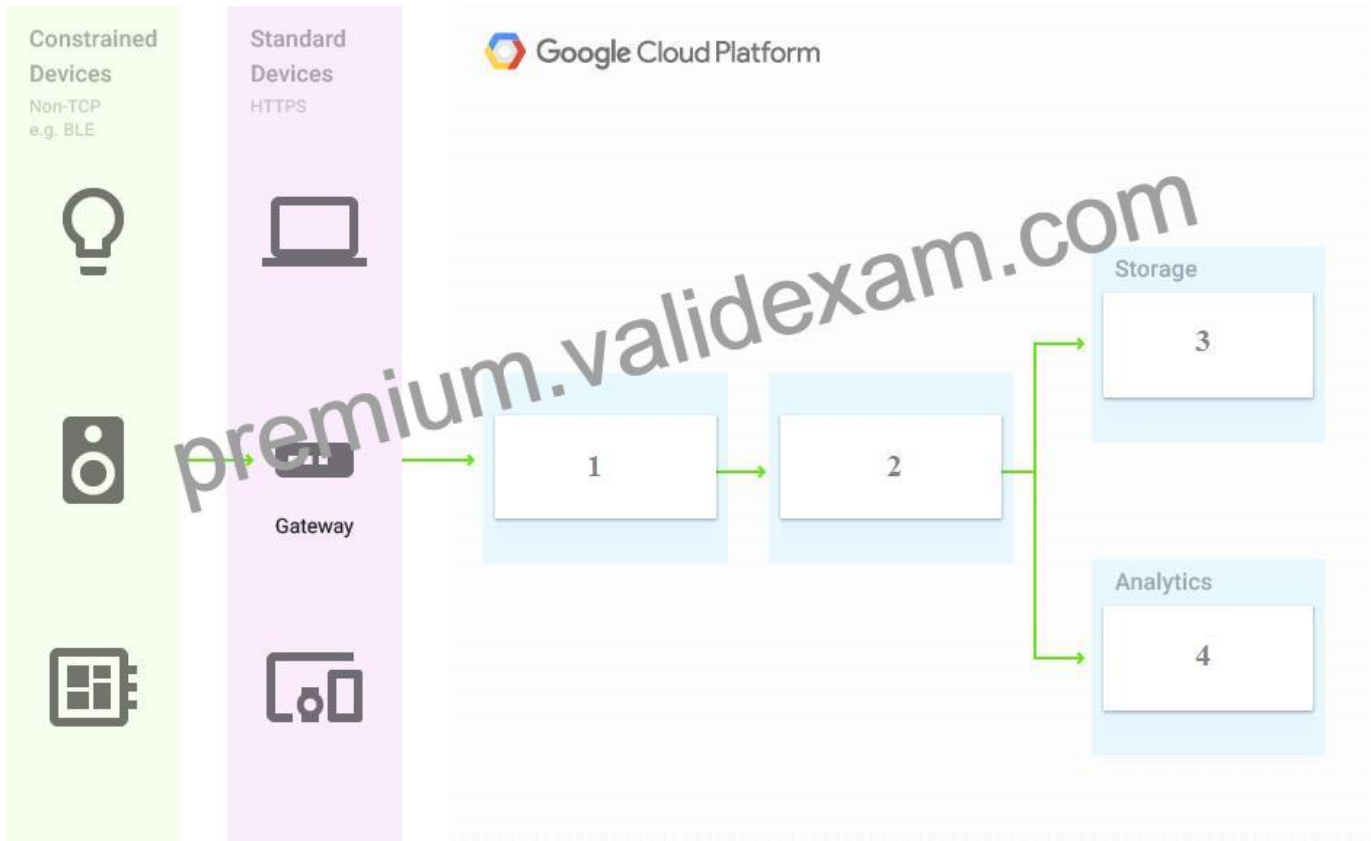
* Cloud Pub/Sub, Cloud Dataflow, Cloud Datastore, BigQuery
* Firebase Messages, Cloud Pub/Sub, Cloud Spanner, BigQuery
* Cloud Pub/Sub, Cloud Storage, BigQuery, Cloud Bigtable
* Cloud Pub/Sub, Cloud Dataflow, Cloud Bigtable, BigQuery

Reference:

https://cloud.google.com/solutions/correlating-time-series-dataflow

QUESTION 53

You have an App Engine application serving as your front-end. It’s going to publish messages to Pub/Sub. The Pub/Sub API hasn’t been enabled yet. What is the fastest way to enable the API?

* Use a service account to auto-enable the API.
* Enable the API in the Console.
* Application’s in App Engine don’t require external APIs to be enabled.
* The API will be enabled the first time the code attempts to access Pub/Sub.

QUESTION 54

You create a Deployment with 2 replicas in a Google Kubernetes Engine cluster that has a single preemptible node pool. After a few minutes, you use kubectlto examine the status of your Pod and observe that one of them is still in Pendingstatus:

```
$ kubectl get pods -l app=myapp
NAME                              READY   STATUS    RESTART   AGE
myapp-deployment-58ddbbb995-1p86m 0/1     Pending   0         9m
myapp-deployment-58ddbbb995-qjpkg 1/1     Running   0         9m
```

What is the most likely cause?

* The pending Pod’s resource requests are too large to fit on a single node of the cluster.

* Too many Pods are already running in the cluster, and there are not enough resources left to schedule the pending Pod.

* The node pool is configured with a service account that does not have permission to pull the container image used by the pending Pod.

* The pending Pod was originally scheduled on a node that has been preempted between the creation of the Deployment and your verification of the Pods’ status. It is currently being rescheduled on a new node.

As the node on which pod was scheduled to run was preempted & now this pod is scheduled to run on different preemtible node from the node-pool.

QUESTION 55

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.

QUESTION 56

You have successfully created a development environment in a project for an application. This application uses Compute Engine and Cloud SQL. Now, you need to create a production environment for this application. The security team has forbidden the existence of network routes between these 2 environments, and asks you to follow Google-recommended practices. What should you do?

* Create a new project, enable the Compute Engine and Cloud SQL APIs in that project, and replicate the setup you have created in the development environment.

* Create a new production subnet in the existing VPC and a new production Cloud SQL instance in your existing project, and deploy your application using those resources.

* Create a new project, modify your existing VPC to be a Shared VPC, share that VPC with your new project, and replicate the setup you have in the development environment in that new project, in the Shared VPC.

* Ask the security team to grant you the Project Editor role in an existing production project used by another division of your company. Once they grant you that role, replicate the setup you have in the development environment in that project.

QUESTION 57

While working on a project, an application administrator has been given the responsibility of managing all resources. He wants to delegate the responsibility of managing the existing service accounts to another administrator. He will also be responsible to manage the other service accounts that will be created. Which of the following is the best way to delegate the privileges required to manage all the service accounts?

* Granting iam.serviceAccountUser to the administrator at the project level

* Granting iam.serviceProjectAccountUser to the administrator at the project level

* Granting iam.serviceAccountUser to the administrator at the service account level

* Granting iam.serviceProjectAccountUser to the administrator at the service account level

QUESTION 58

An application requires block storage for file updates. The data is 500 GB and must continuously sustain 100 MiB/s of aggregate read/write operations.

Which storage option is appropriate for this application?
* Amazon S3
* Amazon EFS
* Amazon EBS
* Amazon Glacier

QUESTION 59

You are migrating a production-critical on-premises application that requires 96 vCPUs to perform its task. You want to make sure the application runs in a similar environment on GCP. What should you do?
* When creating the VM, use machine type n1-standard-96.
* When creating the VM, use Intel Skylake as the CPU platform.
* Create the VM using Compute Engine default settings. Use gcloud to modify the running instance to have 96 vCPUs.
* Start the VM using Compute Engine default settings, and adjust as you go based on Rightsizing Recommendations.

QUESTION 60

Users submit requests to a service that takes several minutes to process. A Solutions Architect needs to ensure that these requests are processed at least once, and that the service has the ability to handle large increases in the number of requests.

How should these requirements be met?
* Put the requests into an Amazon SQS queue and configure Amazon EC2 instances to poll the queue
* Publish the message to an Amazon SNS topic that an Amazon EC2 subscriber can receive and process
* Save the requests to an Amazon DynamoDB table with a DynamoDB stream that triggers an Amazon EC2 Spot Instance
* Use Amazon S3 to store the requests and configure an event notification to have Amazon EC2 instances process the new object

QUESTION 61

You need a dynamic way of provisioning VMs on Compute Engine. The exact specifications will be in a dedicated configuration file. You want to follow Google’s recommended practices. Which method should you use?
* Deployment Manager
* Cloud Composer
* Managed Instance Group
* Unmanaged Instance Group

QUESTION 62

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

:

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 63

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?
* – Create Compute Engine resources in us-central1-b.

-Balance the load across both us-central1-a and us-central1-b.
* – Create a Managed Instance Group and specify us-central1-a as the zone.

-Configure the Health Check with a short Health Interval.
* – Create an HTTP(S) Load Balancer.

-Create one or more global forwarding rules to direct traffic to your VMs.
* – Perform regular backups of your application.

-Create a Cloud Monitoring Alert and be notified if your application becomes unavailable.

-Restore from backups when notified.

QUESTION 64

Your management has asked an external auditor to review all the resources in a specific project.

The security team has enabled the Organization Policy called Domain Restricted Sharing on the organization node by specifying only your Cloud Identity domain. You want the auditor to only be able to view, but not modify, the resources in that project. What should you do?
* Ask the auditor for their Google account, and give them the Viewer role on the project.
* Ask the auditor for their Google account, and give them the Security Reviewer role on the project.
* Create a temporary account for the auditor in Cloud Identity, and give that account the Viewer role on the project.
* Create a temporary account for the auditor in Cloud Identity, and give that account the Security Reviewer role on the project.
https://cloud.google.com/iam/docs/roles-audit-logging#scenario_external_auditors

QUESTION 65

You want to configure a solution for archiving data in a Cloud Storage bucket. The solution must be cost- effective. Data with multiple versions should be archived after 30 days. Previous versions are accessed once a month for reporting. This archive data is also occasionally updated at month-end. What should you do?
* Add a bucket lifecycle rule that archives data with newer versions after 30 days to Coldline Storage.
* Add a bucket lifecycle rule that archives data with newer versions after 30 days to Nearline Storage.
* Add a bucket lifecycle rule that archives data from regional storage after 30 days to Coldline Storage.
* Add a bucket lifecycle rule that archives data from regional storage after 30 days to Nearline Storage.
Explanation/Reference: https://cloud.google.com/storage/docs/managing-lifecycles

The Google Associate Cloud Engineer exam targets software engineers who would like to dabble into the Google Cloud. So, you need to be capable of using the command-line interfaces and Google Cloud Console to handle core platform-based tasks. These mostly include monitoring operations and managing enterprise solutions to maintain deployed applications. In addition, after successful writing of the test, candidates obtain the appropriate Google Cloud certification.

Associate-Cloud-Engineer Dumps and Exam Test Engine:

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