ASSOCIATE-CLOUD-ENGINEER^{Q&As}

Associate Cloud Engineer

Pass Google ASSOCIATE-CLOUD-ENGINEER Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.pass4itsure.com/associate-cloud-engineer.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Google Official Exam Center

🔧 **Instant Download** After Purchase

🔧 **100% Money Back** Guarantee

🔧 **365 Days** Free Update

🔧 **800,000+** Satisfied Customers

QUESTION 1

The sales team has a project named Sales Data Digest that has the ID acme-data-digest. You need to set up similar Google Cloud resources for the marketing team but their resources must be organized independently of the sales team. What should you do?

A. Grant the Project Editor role to the Marketing team for acme-data-digest.

B. Create a Project Lien on acme-data-digest and then grant the Project Editor role to the Marketing team.

C. Create another project with the ID acme-marketing-data-digest for the Marketing team and deploy the resources there.

D. Create a new project named Marketing Data Digest and use the ID acme-data-digest. Grant the Project Editor role to the Marketing team.

Correct Answer: C

Answer should be C because the resources for the marketing team should be independent from the Sales team. Resources are tied and separated by projects.

QUESTION 2

You recently received a new Google Cloud project with an attached billing account where you will work. You need to create instances, set firewalls, and store data in Cloud Storage. You want to follow Google-recommended practices. What should you do?

A. Use the gcloud CLI services enable cloudresourcemanager.googleapis.com command to enable all resources.

B. Use the gcloud services enable compute.googleapis.com command to enable Compute Engine and the gcloud services enable storage-api.googleapis.com command to enable the Cloud Storage APIs.

C. Open the Google Cloud console and enable all Google Cloud APIs from the API dashboard.

D. Open the Google Cloud console and run gcloud init --project in a Cloud Shell.

Correct Answer: B

At first you need to enable API

QUESTION 3

You manage three Google Cloud projects with the Cloud Monitoring API enabled. You want to follow Google-recommended practices to visualize CPU and network metrics for all three projects together. What should you do?

A. 1. Create a Cloud Monitoring Dashboard

2.

 Collect metrics and publish them into the Pub/Sub topics

3.

 Add CPU and network Charts (or each of (he three projects

B. 1. Create a Cloud Monitoring Dashboard.

2.

 Select the CPU and Network metrics from the three projects.

3.

 Add CPU and network Charts lot each of the three protects.

C. 1 Create a Service Account and apply roles/viewer on the three projects

2.

 Collect metrics and publish them lo the Cloud Monitoring API

3.

 Add CPU and network Charts for each of the three projects.

D. 1. Create a fourth Google Cloud project 2 Create a Cloud Workspace from the fourth project and add the other three projects

Correct Answer: B

QUESTION 4

You want to configure an SSH connection to a single Compute Engine instance for users in the dev1 group. This instance is the only resource in this particular Google Cloud Platform project that the dev1 users should be able to connect to. What should you do?

A. Set metadata to enable-oslogin=true for the instance. Grant the dev1 group the compute.osLogin role. Direct them to use the Cloud Shell to ssh to that instance.

B. Set metadata to enable-oslogin=true for the instance. Set the service account to no service account for that instance. Direct them to use the Cloud Shell to ssh to that instance.

C. Enable block project wide keys for the instance. Generate an SSH key for each user in the dev1 group. Distribute the keys to dev1 users and direct them to use their third-party tools to connect.

D. Enable block project wide keys for the instance. Generate an SSH key and associate the key with that instance. Distribute the key to dev1 users and direct them to use their third-party tools to connect.

Correct Answer: A

Granting OS Login IAM roles

After you enable OS Login on one or more instances in your project, those VMs accept connections only from user accounts that have the necessary IAM roles in your project or organization.

QUESTION 5

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?

A. Ask the auditor for their Google account, and give them the Viewer role on the project.

B. Ask the auditor for their Google account, and give them the Security Reviewer role on the project.

C. Create a temporary account for the auditor in Cloud Identity, and give that account the Viewer role on the project.

D. Create a temporary account for the auditor in Cloud Identity, and give that account the Security Reviewer role on the project.

Correct Answer: C

oles/viewer Read access to all resources. Get and list access for all resources.

Using primitive roles The following table lists the primitive roles that you can grant to access a project, the description of what the role does, and the permissions bundled within that role. Avoid using primitive roles except when absolutely necessary. These roles are very powerful, and include a large number of permissions across all Google Cloud services. For more details on when you should use primitive roles, see the Identity and Access Management FAQ.

IAM predefined roles are much more granular, and allow you to carefully manage the set of permissions that your users have access to. See Understanding Roles for a list of roles that can be granted at the project level. Creating custom roles can further increase the control you have over user permissions.

https://cloud.google.com/resource-manager/docs/access-control-proj#using_primitive_roles

QUESTION 6

You\'re looking for the IP address of a specific instance that is running in your default zone. Which command and flag(s) could you use to return just the IP address?

A. The gcloud compute list along with the filter and format flags.

B. The gcloud compute list along with the o flag and jsonpath value.

C. The gcloud compute instances list along with the o flag and jsonpath value.

D. The gcloud compute instances list along with the filter and format flags.

Correct Answer: D

QUESTION 7

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 now version. Update the service to use the new deployment.

Correct Answer: A

A is correct answer,

Keyword, Version, traffic splitting, App Engine supports traffic splitting for versions before releasing.

QUESTION 8

You have a Cloud Storage bucket that needs to host static web assets. How do you make the bucket public?

A. Trick question. Don\\'t ever make a bucket public.

B. Check the "make public" box in the UI.

C. Set allUsers to have the Storage Object Viewer role.

D. gsutil make-public gs://bucket-name

Correct Answer: C

QUESTION 9

You\\'re running an n-tier application on Compute Engine with an Apache web server serving up web requests. You want to consolidate all of your logging into Stackdriver. What\\'s the best approach to get the Apache logs into Stackdriver?

A. Create a log sink and export it to Stackdriver.

B. Stackdriver logs application data from all instances by default.

C. Enable Stackdriver monitoring when creating the instance.

D. Install the Stackdriver monitoring and logging agents on the instance.

Correct Answer: D

QUESTION 10

A member of the finance team informed you that one of the projects is using the old billing account. What steps should you take to resolve the problem?

A. Go to the Project page; expand the Billing tile; select the Billing Account option; select the correct billing account and save.

B. Go to the Billing page; view the list of projects; find the project in question and select Change billing account; select the correct billing account and save.

C. Delete the project and recreate it with the correct billing account.

D. Submit a support ticket requesting the change.

Correct Answer: B

QUESTION 11

You are hosting an application on bare-metal servers in your own data center. The application needs access to Cloud Storage. However, security policies prevent the servers hosting the application from having public IP addresses or access to the internet. You want to follow Google-recommended practices to provide the application with access to Cloud Storage. What should you do?

A. 1. Use nslookup to get the IP address for storage.googleapis.com.

2.

 Negotiate with the security team to be able to give a public IP address to the servers.

3.

 Only allow egress traffic from those servers to the IP addresses for storage.googleapis.com.

B. 1. Using Cloud VPN, create a VPN tunnel to a Virtual Private Cloud (VPC) in Google Cloud Platform (GCP).

2.

 In this VPC, create a Compute Engine instance and install the Squid proxy server on this instance.

3.

 Configure your servers to use that instance as a proxy to access Cloud Storage.

C. 1. Use Migrate for Compute Engine (formerly known as Velostrata) to migrate those servers to Compute Engine.

2.

 Create an internal load balancer (ILB) that uses storage.googleapis.com as backend.

3.

 Configure your new instances to use this ILB as proxy.

D. 1. Using Cloud VPN or Interconnect, create a tunnel to a VPC in GCP.

2.

 Use Cloud Router to create a custom route advertisement for 199.36.153.4/30. Announce that network to your on-premises network through the VPN tunnel.

3.

 In your on-premises network, configure your DNS server to resolve *.googleapis.com as a CNAME to
restricted.googleapis.com.

Correct Answer: D

VPN/interconnect without exposing to public

QUESTION 12

You are managing a project for the Business Intelligence (BI) department in your company. A data pipeline ingests data
into BigQuery via streaming. You want the users in the BI department to be able to run the custom SQL queries against
the latest data in BigQuery. What should you do?

A. Create a Data Studio dashboard that uses the related BigQuery tables as a source and give the BI team view access
to the Data Studio dashboard.

B. Create a Service Account for the BI team and distribute a new private key to each member of the BI team.

C. Use Cloud Scheduler to schedule a batch Dataflow job to copy the data from BigQuery to the BI team\\'s internal data
warehouse.

D. Assign the IAM role of BigQuery User to a Google Group that contains the members of the BI team.

Correct Answer: D

When applied to a dataset, this role provides the ability to read the dataset\\'s metadata and list tables in the dataset.
When applied to a project, this role also provides the ability to run jobs, including queries, within the project. A member
with this role can enumerate their own jobs, cancel their own jobs, and enumerate datasets within a project. Additionally,
allows the creation of new datasets within the project; the creator is granted the BigQuery Data Owner role (roles/
bigquery.dataOwner) on these new datasets. https://cloud.google.com/bigquery/docs/access-control

---

QUESTION 13

You are about to deploy a new Enterprise Resource Planning (ERP) system on Google Cloud. The application holds the
full database in-memory for fast data access, and you need to configure the most appropriate resources on Google
Cloud for this application. What should you do?

A. Provision preemptible Compute Engine instances.

B. Provision Compute Engine instances with GPUs attached.

C. Provision Compute Engine instances with local SSDs attached.

D. Provision Compute Engine instances with M1 machine type.

Correct Answer: D

M1 machine series Medium in-memory databases such as SAP HANA Tasks that require intensive use of memory with
higher memory-to-vCPU ratios than the general- purpose high-memory machine types. In-memory databases and in-

memory analytics, business warehousing (BW) workloads, genomics analysis, SQL analysis services.

Microsoft SQL Server and similar databases.

https://cloud.google.com/compute/docs/machine-types https://cloud.google.com/compute/docs/machine-types#:~:text=databases%20such%20as-,SAP%20HANA,-In%2Dmemory%20databases
https://www.sap.com/india/products/

hana.html#:~:text=is%20SAP%20HANA-,in%2Dmemory,-database%3F

Reference: https://cloud.google.com/compute/docs/disks/local-ssd

---

QUESTION 14

You have deployed multiple Linux instances on Compute Engine. You plan on adding more instances in the coming weeks. You want to be able to access all of these instances through your SSH client over the internet without having to configure specific access on the existing and new instances. You do not want the Compute Engine instances to have a public IP. What should you do?

A. Configure Cloud Identity-Aware Proxy for HTTPS resources.

B. Configure Cloud Identity-Aware Proxy for SSH and TCP resources

C. Create an SSH keypair and store the public key as a project-wide SSH Key.

D. Create an SSH keypair and store the private key as a project-wide SSH Key.

Correct Answer: B

https://cloud.google.com/iap/docs/using-tcp-forwarding

QUESTION 15

You have files in a Cloud Storage bucket that you need to share with your suppliers. You want to restrict the time that the files are available to your suppliers to 1 hour. You want to follow Google recommended practices. What should you do?

A. Create a service account with just the permissions to access files in the bucket. Create a JSON key for the service account. Execute the command gsutil signurl -m 1h gs:///*.

B. Create a service account with just the permissions to access files in the bucket. Create a JSON key for the service account. Execute the command gsutil signurl -d 1h gs:///**.

C. Create a service account with just the permissions to access files in the bucket. Create a JSON key for the service account. Execute the command gsutil signurl -p 60m gs:///.

D. Create a JSON key for the Default Compute Engine Service Account. Execute the command gsutil signurl -t 60m gs:///***

Correct Answer: B

Explanation: This command correctly specifies the duration that the signed url should be valid for by using the -d flag. The default is 1 hour so omitting the -d flag would have also resulted in the same outcome. Times may be specified with

no suffix (default hours), or with s = seconds, m = minutes, h = hours, d = days. The max duration allowed is 7d.Ref:
https://cloud.google.com/storage/docs/gsutil/commands/signurl

| Latest ASSOCIATE-CLOUD-ENGINEER Dumps | ASSOCIATE-CLOUD-ENGINEER Practice Test | ASSOCIATE-CLOUD-ENGINEER Study Guide |