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

QUESTION 1

You want to verify the IAM users and roles assigned within a GCP project named my-project. What should you do?

A. Run gcloud iam roles list. Review the output section.

B. Run gcloud iam service-accounts list. Review the output section.

C. Navigate to the project and then to the IAM section in the GCP Console. Review the members and roles.

D. Navigate to the project and then to the Roles section in the GCP Console. Review the roles and status.

Correct Answer: C

C is the only logical answers. If you go IAM and Admin > IAM: You can see Principals and Roles. Users, groups, service accounts are Principals

QUESTION 2

You want to deploy an application on Cloud Run that processes messages from a Cloud Pub/Sub topic. You want to follow Google-recommended practices. What should you do?

A. 1. Create a Cloud Function that uses a Cloud Pub/Sub trigger on that topic.

2. Call your application on Cloud Run from the Cloud Function for every message.

B. 1. Grant the Pub/Sub Subscriber role to the service account used by Cloud Run.

2.

Create a Cloud Pub/Sub subscription for that topic.

3.

Make your application pull messages from that subscription.

C. 1. Create a service account.

2.

Give the Cloud Run Invoker role to that service account for your Cloud Run application.

3.

Create a Cloud Pub/Sub subscription that uses that service account and uses your Cloud Run application as the push endpoint.

D. 1. Deploy your application on Cloud Run on GKE with the connectivity set to Internal.

2.

Create a Cloud Pub/Sub subscription for that topic.

3.

In the same Google Kubernetes Engine cluster as your application, deploy a container that takes the messages and sends them to your application.

Correct Answer: C

C. is the correct answer: Create a service account. 2. Give the Cloud Run Invoker role to that service account for your Cloud Run application. 3. Create a Cloud Pub/Sub subscription that uses that service account and uses your Cloud Run application as the push endpoint

QUESTION 3

Your organization has strict requirements to control access to Google Cloud projects. You need to enable your Site Reliability Engineers (SREs) to approve requests from the Google Cloud support team when an SRE opens a support case. You want to follow Google-recommended practices. What should you do?

A. Add your SREs to roles/iam.roleAdmin role.

B. Add your SREs to roles/accessapproval approver role.

C. Add your SREs to a group and then add this group to roles/iam roleAdmin role.

D. Add your SREs to a group and then add this group to roles/accessapproval approver role.

Correct Answer: D

Add your SREs to a group and then add this group to roles/accessapproval approver role. -Google recommendation.

QUESTION 4

For analysis purposes, you need to send all the logs from all of your Compute Engine instances to a BigQuery dataset called platform-logs. You have already installed the Stackdriver Logging agent on all the instances. You want to minimize cost. What should you do?

A. 1. Give the BigQuery Data Editor role on the platform-logs dataset to the service accounts used by your instances.

2. Update your instances\\' metadata to add the following value: logs-destination: bq://platform-logs.

B. 1. In Stackdriver Logging, create a logs export with a Cloud Pub/Sub topic called logs as a sink.

2.

Create a Cloud Function that is triggered by messages in the logs topic.

3.

Configure that Cloud Function to drop logs that are not from Compute Engine and to insert Compute Engine logs in the platform-logs dataset.

C. 1. In Stackdriver Logging, create a filter to view only Compute Engine logs.

2.

Click Create Export.

3.

Choose BigQuery as Sink Service, and the platform-logs dataset as Sink Destination.

D. 1. Create a Cloud Function that has the BigQuery User role on the platform-logs dataset.

2.

Configure this Cloud Function to create a BigQuery Job that executes this query:

INSERT INTO dataset.platform-logs (timestamp, log)

SELECT timestamp, log FROM compute.logs

WHERE timestamp > DATE_SUB(CURRENT_DATE(), INTERVAL 1 DAY)

3.

Use Cloud Scheduler to trigger this Cloud Function once a day.

Correct Answer: C

C. is correct, Sinks control how Cloud Logging routes logs. Using sinks, you can route some or all of your logs to supported destinations.

QUESTION 5

You need to run an important query in BigQuery but expect it to return a lot of records. You want to find out how much it will cost to run the query. You are using on-demand pricing. What should you do?

A. Arrange to switch to Flat-Rate pricing for this query, then move back to on-demand.

B. Use the command line to run a dry run query to estimate the number of bytes read. Then convert that bytes estimate to dollars using the Pricing Calculator.

C. Use the command line to run a dry run query to estimate the number of bytes returned. Then convert that bytes estimate to dollars using the Pricing Calculator.

D. Run a select count (*) to get an idea of how many records your query will look through. Then convert that number of rows to dollars using the Pricing Calculator.

Correct Answer: B

Reference: https://cloud.google.com/bigquery/docs/estimate-costs

Under on-demand pricing, BigQuery charges for queries by using one metric: the number of bytes processed (also referred to as bytes read). You are charged for the number of bytes processed whether the data is stored in BigQuery or in an external data source such as Cloud Storage, Drive, or Cloud Bigtable. On-demand pricing is based solely on usage.

QUESTION 6

You are given a project with a single virtual private cloud (VPC) and a single subnetwork in the us-central1 region. There is a Compute Engine instance hosting an application in this subnetwork. You need to deploy a new instance in the same project in the europe-west1 region. This new instance needs access to the application. You want to follow Google-recommended practices. What should you do?

A. 1. Create a subnetwork in the same VPC, in europe-west1.

2. Create the new instance in the new subnetwork and use the first instance\\'s private address as the endpoint.

B. 1. Create a VPC and a subnetwork in europe-west1.

2.

Expose the application with an internal load balancer.

3.

Create the new instance in the new subnetwork and use the load balancer\\'s address as the endpoint.

C. 1. Create a subnetwork in the same VPC, in europe-west1.

2.

Use Cloud VPN to connect the two subnetworks.

3.

Create the new instance in the new subnetwork and use the first instance\\'s private address as the endpoint.

D. 1. Create a VPC and a subnetwork in europe-west1.

2.

Peer the 2 VPCs.

3.

Create the new instance in the new subnetwork and use the first instance\\'s private address as the endpoint.

Correct Answer: A

VPC allows you to spawn multiple subnets in different zones. Routing is handled automatically (because Routers are created automatically).

"use the first instance\\'s private address as the endpoint" means that this new instance will be accessing the app via first intance\\'s private IP (so there should be some routing rules created). Question says: "This new instance needs access to the application."

QUESTION 7

You want to find out when users were added to Cloud Spanner Identity Access Management (IAM) roles on your Google Cloud Platform (GCP) project. What should you do in the GCP Console?

A. Open the Cloud Spanner console to review configurations.

B. Open the IAM and admin console to review IAM policies for Cloud Spanner roles.

C. Go to the Stackdriver Monitoring console and review information for Cloud Spanner.

D. Go to the Stackdriver Logging console, review admin activity logs, and filter them for Cloud Spanner IAM roles.

Correct Answer: D

Activity logs captures the time when the users were given the IAM roles for Cloud Spanner

QUESTION 8

Your application development team has created Docker images for an application that will be deployed on Google Cloud. Your team does not want to manage the infrastructure associated with this application. You need to ensure that the application can scale automatically as it gains popularity. What should you do?

A. Create an instance template with the container image, and deploy a Managed Instance Group with Autoscaling.

B. Upload Docker images to Artifact Registry, and deploy the application on Google Kubernetes Engine using Standard mode.

C. Upload Docker images to the Cloud Storage, and deploy the application on Google Kubernetes Engine using Standard mode.

D. Upload Docker images to Artifact Registry, and deploy the application on Cloud Run.

Correct Answer: D

QUESTION 9

You have an application that looks for its licensing server on the IP 10.0.3.21. You need to deploy the licensing server on Compute Engine. You do not want to change the configuration of the application and want the application to be able to reach the licensing server. What should you do?

A. Reserve the IP 10.0.3.21 as a static internal IP address using gcloud and assign it to the licensing server.

B. Reserve the IP 10.0.3.21 as a static public IP address using gcloud and assign it to the licensing server.

C. Use the IP 10.0.3.21 as a custom ephemeral IP address and assign it to the licensing server.

D. Start the licensing server with an automatic ephemeral IP address, and then promote it to a static internal IP address.

Correct Answer: A

IP 10.0.3.21 is internal by default, and to ensure that it will be static non-changing it should be selected as static internal ip address.

QUESTION 10

Your company uses BigQuery for data warehousing. Over time, many different business units in your company have

created 1000+ datasets across hundreds of projects. Your CIO wants you to examine all datasets to find tables that contain an employee_ssn column. You want to minimize effort in performing this task. What should you do?

A. Go to Data Catalog and search for employee_ssn in the search box.

B. Write a shell script that uses the bq command line tool to loop through all the projects in your organization.

C. Write a script that loops through all the projects in your organization and runs a query on INFORMATION_SCHEMA.COLUMNS view to find the employee_ssn column.

D. Write a Cloud Dataflow job that loops through all the projects in your organization and runs a query on INFORMATION_SCHEMA.COLUMNS view to find employee_ssn column.

Correct Answer: A

A is the correct answer, Data Catalog can be used to search the column with keyword:value pair,

Filter your search by adding a keyword:value to your search terms in the search box:

Keyword Description name: Match data asset name ***column: Match column name or nested column name description: Match table description

---

QUESTION 11

Your developers have created an application that needs to be able to make calls to Cloud Storage and BigQuery. The code is going to run inside a container and will run on Kubernetes Engine and on- premises. What\\'s the best way for them to authenticate to the Google Cloud services?

A. Create a service account, grant it the least viable privileges to the required services, generate and download a key. Use the key to authenticate inside the application.

B. Use the default service account for App Engine which already has the required permissions.

C. Use the default service account for Compute Engine which already has the required permissions.

D. Create a service account, with editor permissions, generate and download a key. Use the key to authenticate inside the application.

Correct Answer: A

QUESTION 12

You deployed an App Engine application using gcloud app deploy, but it did not deploy to the intended project. You want to find out why this happened and where the application deployed. What should you do?

A. Check the app.yaml file for your application and check project settings.

B. Check the web-application.xml file for your application and check project settings.

C. Go to Deployment Manager and review settings for deployment of applications.

D. Go to Cloud Shell and run gcloud config list to review the Google Cloud configuration used for deployment.

Correct Answer: D

check project setting by gcloud config list

QUESTION 13

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 14

Your boss has asked you to set up the CFO as a user inside Google Cloud. Before walking away, she said, "And obviously you know which role to assign him." Which role was she likely talking about, and why is it obvious? (Select one role and one reason)

A. Reason: This role will ensure the CFO has full access to the project.

B. Role: Billing Account Viewer

C. Reason: This role will ensure the CFO has access to view the spending data;

D. Role: Project Owner

E. Role: Billing Account Admin

F. Reason:: This role will ensure the CFO can perform any billing task that they may need.

Correct Answer: BC

QUESTION 15

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.

Correct Answer: A

A for sure

B - this allows internal communicaiton, but does nothing to limit public traffic C - deny all is nice, but it\\'s for egress -- we\\'re looking for ingress D - this is way to invasive and doesn\\'t explicitly address the issue of preventing public internet traffic from reaching your instance -- if it does, someone let me know how.