



Free Questions for Professional-Cloud-Architect

Shared by Mercer on 29-01-2024

For More Free Questions and Preparation Resources

[Check the Links on Last Page](#)



Question 1

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

Professional-Cloud-Architect Mix Questions IN THIS CASE STUDY

You are migrating third-party applications from optimized on-premises virtual machines to Google Cloud. You are unsure about the optimum CPU and memory options. The application have a consistent usage patterns across multiple weeks. You want to optimize resource usage for the lowest cost. What should you do?

Options:

- A- Create a Compute engine instance with CPU and Memory options similar to your application's current on-premises virtual machine. Install the cloud monitoring agent, and deploy the third party application. Run a load with normal traffic levels on third party application and follow the Rightsizing Recommendations in the Cloud Console
- B- Create an App Engine flexible environment, and deploy the third party application using a Docker file and a custom runtime. Set CPU and memory options similar to your application's current on-premises virtual machine in the app.yaml file.
- C- Create an instance template with the smallest available machine type, and use an image of the third party application taken from the current on-premises virtual machine. Create a managed instance group that uses average CPU to autoscale the number of instances in the group. Modify the average CPU utilization threshold to optimize the number of instances running.
- D- Create multiple Compute Engine instances with varying CPU and memory options. Install the cloud monitoring agent and deploy the third-party application on each of them. Run a load test with high traffic levels on the application and use the results to determine the optimal settings.

Answer:

A

Explanation:

Create a Compute engine instance with CPU and Memory options similar to your application's current on-premises virtual machine. Install the cloud monitoring agent, and deploy the third party application. Run a load with normal traffic levels on third party application and follow the Rightsizing Recommendations in the Cloud Console

<https://cloud.google.com/migrate/compute-engine/docs/4.9/concepts/planning-a-migration/cloud-i>

[nstance-rightsizing?hl=en](#)

Question 2

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

Professional-Cloud-Architect Mix Questions IN THIS CASE STUDY

You need to deploy a stateful workload on Google Cloud. The workload can scale horizontally, but each instance needs to read and write to the same POSIX filesystem. At high load, the stateful workload needs to support up to 100 MB/s of writes. What should you do?

Options:

- A- Use a persistent disk for each instance.
- B- Use a regional persistent disk for each instance.
- C- Create a Cloud Filestore instance and mount it in each instance.
- D- Create a Cloud Storage bucket and mount it in each instance using gcsfuse.

Answer:

C

Explanation:

<https://cloud.google.com/storage/docs/gcs-fuse#notes>

Cloud Filestore: Cloud Filestore is a scalable and highly available shared file service fully managed by Google. Cloud Filestore provides persistent storage ideal for shared workloads. It is best suited for enterprise applications requiring persistent, durable, shared storage which is accessed by NFS or requires a POSIX compliant file system.

Question 3

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

Professional-Cloud-Architect Mix Questions IN THIS CASE STUDY

Your company pushes batches of sensitive transaction data from its application server VMs to Cloud Pub/Sub for processing and storage. What is the Google-recommended way for your application to authenticate to the required Google Cloud services?

Options:

- A- Ensure that VM service accounts are granted the appropriate Cloud Pub/Sub IAM roles.
- B- Ensure that VM service accounts do not have access to Cloud Pub/Sub, and use VM access scopes to grant the appropriate Cloud Pub/Sub IAM roles.
- C- Generate an OAuth2 access token for accessing Cloud Pub/Sub, encrypt it, and store it in Cloud Storage for access from each VM.
- D- Create a gateway to Cloud Pub/Sub using a Cloud Function, and grant the Cloud Function service account the appropriate Cloud Pub/Sub IAM roles.

Answer:

A

Question 4

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

Professional-Cloud-Architect Mix Questions IN THIS CASE STUDY

Your company captures all web traffic data in Google Analytics 360 and stores it in BigQuery. Each country has its own dataset. Each dataset has multiple tables. You want analysts from each country

to be able to see and query only the data for their respective countries.

How should you configure the access rights?

Options:

- A-** Create a group per country. Add analysts to their respective country-groups. Create a single group 'all_analysts', and add all country-groups as members. Grant the 'all-analysis' group the IAM role of BigQuery jobUser. Share the appropriate dataset with view access with each respective analyst country-group.
- B-** Create a group per country. Add analysts to their respective country-groups. Create a single group 'all_analysts', and add all country-groups as members. Grant the 'all-analysis' group the IAM role of BigQuery jobUser. Share the appropriate tables with view access with each respective analyst countrygroup.
- C-** Create a group per country. Add analysts to their respective country-groups. Create a single group 'all_analysts', and add all country-groups as members. Grant the 'all-analysis' group the IAM role of BigQuery dataViewer. Share the appropriate dataset with view access with each respective analyst country-group.
- D-** Create a group per country. Add analysts to their respective country-groups. Create a single group 'all_analysts', and add all country-groups as members. Grant the 'all-analysis' group the IAM role of BigQuery dataViewer. Share the appropriate table with view access with each respective analyst countrygroup.

Answer:

A

Question 5

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

Professional-Cloud-Architect Mix Questions IN THIS CASE STUDY

Your company is developing a new application that will allow globally distributed users to upload pictures and share them with other selected users. The application will support millions of concurrent users. You want to allow developers to focus on just building code without having to

create and maintain the underlying infrastructure. Which service should you use to deploy the application?

Options:

- A- App Engine
- B- Cloud Endpoints
- C- Compute Engine
- D- Google Kubernetes Engine

Answer:

A

Explanation:

<https://cloud.google.com/appengine/docs/standard/go/how-requests-are-handled>

Question 6

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

Professional-Cloud-Architect Mix Questions IN THIS CASE STUDY

Your company has an enterprise application running on Compute Engine that requires high availability and high performance. The application has been deployed on two instances in two zones in the same region in active passive mode. The application writes data to a persistent disk in the case of a single zone outage that data should be immediately made available to the other instance in the other zone. You want to maximize performance while minimizing downtime and data loss. What should you do?

Options:

- A- 1. Attach a persistent SSD disk to the first instance
- 2. Create a snapshot every hour
- 3. In case of a zone outage, recreate a persistent SSD disk in the second instance where data is coming from the created snapshot

- B- 1 Create a Cloud Storage bucket
- 2. Mount the bucket into the first instance with gcs-fuse
- 3. In case of a zone outage, mount the Cloud Storage bucket to the second instance with gcs-fuse
- C- 1 Attach a local SSD to the first instance disk
- 2. Execute an rsync command every hour where the target is a persistent SSD disk attached to the second instance
- 3. In case of a zone outage, use the second instance
- D- 1. Attach a regional SSD persistent disk to the first instance
- 2. In case of a zone outage, force-attach the disk to the other instance

Answer:

D

Question 7

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

Professional-Cloud-Architect Mix Questions IN THIS CASE STUDY

You are working with a data warehousing team that performs data analysis. The team needs to process data from external partners, but the data contains personally identifiable information (PII). You need to process and store the data without storing any of the PII data

a. What should you do?

Options:

- A- Create a Dataflow pipeline to retrieve the data from the external sources. As part of the pipeline use the Cloud Data Loss Prevention (Cloud DLP) API to remove any PII data Store the result in BigQuery
- B- Create a Dataflow pipeline to retrieve the data from the external sources. As part of the pipeline store all non-PII data in BigQuery and store all PII data in a Cloud Storage bucket that has a retention policy set.
- C- Ask the external partners to upload data on Cloud Storage Configure Bucket Lock for the bucket Create a Dataflow pipeline to read the data from the bucket As part of the pipeline, use the Cloud Data Loss Prevention (Cloud DLP) API to remove any PII data Store the result in BigQuery
- D- Ask the external partners to import all data in your BigQuery dataset Create a dataflow pipeline to copy the data into a new table As part of the Dataflow bucket skip all data in columns

that have PII data

Answer:

A

Explanation:

Create a Dataflow pipeline to retrieve the data from the external sources, he did not specify the way he is going to create it, it might be a pub/sub or external table or whatever.

Question 8

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

Professional-Cloud-Architect Mix Questions IN THIS CASE STUDY

You have a Python web application with many dependencies that requires 0.1 CPU cores and 128 MB of memory to operate in production. You want to monitor and maximize machine utilization. You also to reliably deploy new versions of the application. Which set of steps should you take?

Options:

A- Perform the following:

- 1) Create a managed instance group with f1-micro type machines.
- 2) Use a startup script to clone the repository, check out the production branch, install the dependencies, and start the Python app.
- 3) Restart the instances to automatically deploy new production releases.

B- Perform the following:

- 1) Create a managed instance group with n1-standard-1 type machines.
- 2) Build a Compute Engine image from the production branch that contains all of the dependencies and automatically starts the Python app.
- 3) Rebuild the Compute Engine image, and update the instance template to deploy new production releases.

C- Perform the following:

- 1) Create a Kubernetes Engine cluster with n1-standard-1 type machines.

- 2) Build a Docker image from the production branch with all of the dependencies, and tag it with the
- 3) Create a Kubernetes Deployment with the imagePullPolicy set to "IfNotPresent" in the staging namespace, and then promote it to the production namespace after testing.

D- Perform the following:

- 1) Create a Kubernetes Engine (GKE) cluster with n1-standard-4 type machines.
- 2) Build a Docker image from the master branch with all of the dependencies, and tag it with "latest".
- 3) Create a Kubernetes Deployment in the default namespace with the imagePullPolicy set to "Always".

Restart the pods to automatically deploy new production releases.

Answer:

D

Explanation:

<https://cloud.google.com/compute/docs/instance-templates>

Question 9

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

Professional-Cloud-Architect Mix Questions IN THIS CASE STUDY

Your company operates nationally and plans to use GCP for multiple batch workloads, including some that are not time-critical. You also need to use GCP services that are HIPAA-certified and manage service costs.

How should you design to meet Google best practices?

Options:

- A- Provisioning preemptible VMs to reduce cost. Discontinue use of all GCP services and APIs that are not HIPAA-compliant.
- B- Provisioning preemptible VMs to reduce cost. Disable and then discontinue use of all GCP and APIs that are not HIPAA-compliant.

- C- Provision standard VMs in the same region to reduce cost. Discontinue use of all GCP services and APIs that are not HIPAA-compliant.
- D- Provision standard VMs to the same region to reduce cost. Disable and then discontinue use of all GCP services and APIs that are not HIPAA-compliant.

Answer:

B

Explanation:

<https://cloud.google.com/security/compliance/hipaa/>

P2P
exams

P2P
exams

To Get Premium Files for Professional-Cloud-Architect Visit

<https://www.p2pexams.com/products/professional-cloud-architect>

For More Free Questions Visit

<https://www.p2pexams.com/google/pdf/professional-cloud-architect>

20%
DISCOUNT

P2P
exams