



Free Questions for Professional-Cloud-Developer by braindumpscollection

Shared by Leon on 15-04-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You are developing an application using different microservices that should remain internal to the cluster. You want to be able to configure each microservice with a specific number of replicas. You also want to be able to address a specific microservice from any other microservice in a uniform way, regardless of the number of replicas the microservice scales to. You need to implement this solution on Google Kubernetes Engine. What should you do?

Options:

- A-** Deploy each microservice as a Deployment. Expose the Deployment in the cluster using a Service, and use the Service DNS name to address it from other microservices within the cluster.
- B-** Deploy each microservice as a Deployment. Expose the Deployment in the cluster using an Ingress, and use the Ingress IP address to address the Deployment from other microservices within the cluster.
- C-** Deploy each microservice as a Pod. Expose the Pod in the cluster using a Service, and use the Service DNS name to address the microservice from other microservices within the cluster.
- D-** Deploy each microservice as a Pod. Expose the Pod in the cluster using an Ingress, and use the Ingress IP address name to address the Pod from other microservices within the cluster.

Answer:

A

Question 2

Question Type: MultipleChoice

You are developing an application that will store and access sensitive unstructured data objects in a Cloud Storage bucket. To comply with regulatory requirements, you need to ensure that all data objects are available for at least 7 years after their initial creation. Objects created more than 3 years ago are accessed very infrequently (less than once a year). You need to configure object storage while ensuring that storage cost is optimized. What should you do? (Choose two.)

Options:

- A-** Set a retention policy on the bucket with a period of 7 years.
- B-** Use IAM Conditions to provide access to objects 7 years after the object creation date.
- C-** Enable Object Versioning to prevent objects from being accidentally deleted for 7 years after object creation.
- D-** Create an object lifecycle policy on the bucket that moves objects from Standard Storage to Archive Storage after 3 years.
- E-** Implement a Cloud Function that checks the age of each object in the bucket and moves the objects older than 3 years to a second bucket with the Archive Storage class. Use Cloud Scheduler to trigger the Cloud Function on a daily schedule.

Answer:

A, D

Explanation:

<https://cloud.google.com/storage/docs/bucket-lock>

This page discusses the Bucket Lock feature, which allows you to configure a data retention policy for a Cloud Storage bucket that governs how long objects in the bucket must be retained. The feature also allows you to lock the data retention policy, permanently preventing the policy from being reduced or removed.

<https://cloud.google.com/storage/docs/storage-classes#archive>

Archive storage is the lowest-cost, highly durable storage service for data archiving, online backup, and disaster recovery. Unlike the 'coldest' storage services offered by other Cloud providers, your data is available within milliseconds, not hours or days.

Archive storage is the best choice for data that you plan to access less than once a year.

Question 3

Question Type: MultipleChoice

You are planning to deploy hundreds of microservices in your Google Kubernetes Engine (GKE) cluster. How should you secure communication between the microservices on GKE using a managed service?

Options:

- A- Use global HTTP(S) Load Balancing with managed SSL certificates to protect your services
- B- Deploy open source Istio in your GKE cluster, and enable mTLS in your Service Mesh
- C- Install cert-manager on GKE to automatically renew the SSL certificates.
- D- Install Anthos Service Mesh, and enable mTLS in your Service Mesh.

Answer:

D

Explanation:

https://cloud.google.com/service-mesh/docs/overview#security_benefits

- Ensures encryption in transit. Using mTLS for authentication also ensures that all TCP communications are encrypted in transit.

Question 4

Question Type: MultipleChoice

You have an application deployed in Google Kubernetes Engine (GKE). You need to update the application to make authorized requests to Google Cloud managed services. You want this to be a one-time setup, and you need to follow security best practices of auto-rotating your security keys and storing them in an encrypted store. You already created a service account with appropriate access to the Google Cloud service. What should you do next?

Options:

- A-** Assign the Google Cloud service account to your GKE Pod using Workload Identity.
- B-** Export the Google Cloud service account, and share it with the Pod as a Kubernetes Secret.
- C-** Export the Google Cloud service account, and embed it in the source code of the application.
- D-** Export the Google Cloud service account, and upload it to HashiCorp Vault to generate a dynamic service account for your application.

Answer:

A

Explanation:

<https://cloud.google.com/kubernetes-engine/docs/concepts/workload-identity>

Applications running on GKE might need access to Google Cloud APIs such as Compute Engine API, BigQuery Storage API, or Machine Learning APIs.

Workload Identity allows a Kubernetes service account in your GKE cluster to act as an IAM service account. Pods that use the configured Kubernetes service account automatically authenticate as the IAM service account when accessing Google Cloud APIs. Using Workload Identity allows you to assign distinct, fine-grained identities and authorization for each application in your cluster.

Question 5

Question Type: MultipleChoice

You are writing from a Go application to a Cloud Spanner database. You want to optimize your application's performance using Google-recommended best practices. What should you do?

Options:

- A-** Write to Cloud Spanner using Cloud Client Libraries.
- B-** Write to Cloud Spanner using Google API Client Libraries
- C-** Write to Cloud Spanner using a custom gRPC client library.

D- Write to Cloud Spanner using a third-party HTTP client library.

Answer:

A

Explanation:

<https://cloud.google.com/apis/docs/cloud-client-libraries>

"Cloud Client Libraries are the recommended option for accessing Cloud APIs programmatically, where available. Cloud Client Libraries use the latest client library models"

<https://cloud.google.com/apis/docs/client-libraries-explained>

<https://cloud.google.com/go/docs/reference>

Question 6

Question Type: MultipleChoice

You are building a highly available and globally accessible application that will serve static content to users. You need to configure the storage and serving components. You want to minimize management overhead and latency while maximizing reliability for users. What

should you do?

Options:

- A-** 1) Create a managed instance group. Replicate the static content across the virtual machines (VMs)
- 2) Create an external HTTP(S) load balancer.
- 3) Enable Cloud CDN, and send traffic to the managed instance group.
- B-** 1) Create an unmanaged instance group. Replicate the static content across the VMs.
- 2) Create an external HTTP(S) load balancer
- 3) Enable Cloud CDN, and send traffic to the unmanaged instance group.
- C-** 1) Create a Standard storage class, regional Cloud Storage bucket. Put the static content in the bucket
- 2) Reserve an external IP address, and create an external HTTP(S) load balancer
- 3) Enable Cloud CDN, and send traffic to your backend bucket
- D-** 1) Create a Standard storage class, multi-regional Cloud Storage bucket. Put the static content in the bucket.
- 2) Reserve an external IP address, and create an external HTTP(S) load balancer.
- 3) Enable Cloud CDN, and send traffic to your backend bucket.

Answer:

D

Question 7

Question Type: MultipleChoice

You are running a web application on Google Kubernetes Engine that you inherited. You want to determine whether the application is using libraries with known vulnerabilities or is vulnerable to XSS attacks. Which service should you use?

Options:

- A- Google Cloud Armor
- B- Debugger
- C- Web Security Scanner
- D- Error Reporting

Answer:

C

Explanation:

<https://cloud.google.com/security-command-center/docs/concepts-web-security-scanner-overview>

Web Security Scanner identifies security vulnerabilities in your App Engine, Google Kubernetes Engine (GKE), and Compute Engine web applications. It crawls your application, following all links within the scope of your starting URLs, and attempts to exercise as many user inputs and event handlers as possible.

Question 8

Question Type: MultipleChoice

You are developing an application that will handle requests from end users. You need to secure a Cloud Function called by the application to allow authorized end users to authenticate to the function via the application while restricting access to unauthorized users. You will integrate Google Sign-In as part of the solution and want to follow Google-recommended best practices. What should you do?

Options:

- A-** Deploy from a source code repository and grant users the roles/cloudfunctions.viewer role.
- B-** Deploy from a source code repository and grant users the roles/cloudfunctions.invoker role
- C-** Deploy from your local machine using gcloud and grant users the roles/cloudfunctions.admin role
- D-** Deploy from your local machine using gcloud and grant users the roles/cloudfunctions.developer role

Answer:

C

Question 9

Question Type: MultipleChoice

You are designing a chat room application that will host multiple rooms and retain the message history for each room. You have selected Firestore as your database. How should you represent the data in Firestore?

Options:

- A-** Create a collection for the rooms. For each room, create a document that lists the contents of the messages
- B-** Create a collection for the rooms. For each room, create a collection that contains a document for each message
- C-** Create a collection for the rooms. For each room, create a document that contains a collection for documents, each of which contains a message.
- D-** Create a collection for the rooms, and create a document for each room. Create a separate collection for messages, with one document per message. Each room's document contains a list of references to the messages.

Answer:

C

Explanation:

<https://firebase.google.com/docs/firestore/data-model#hierarchical-data>

To Get Premium Files for Professional-Cloud-Developer Visit

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

For More Free Questions Visit

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

